



US 20200258181A1

(19) **United States**

(12) **Patent Application Publication**
SONG et al.

(10) **Pub. No.: US 2020/0258181 A1**

(43) **Pub. Date: Aug. 13, 2020**

(54) **INTELLIGENT REPORT WRITER**

Publication Classification

(71) Applicants: **Yuh-shen SONG**, Porter Ranch, CA (US); **Catherine LEW**, Porter Ranch, CA (US); **Alexander SONG**, Porter Ranch, CA (US); **Victoria SONG**, Porter Ranch, CA (US)

(51) **Int. Cl.**
G06Q 50/26 (2006.01)
G06Q 10/10 (2006.01)
G06Q 20/40 (2006.01)
G06Q 30/00 (2006.01)
G06F 16/22 (2006.01)
G06F 16/28 (2006.01)

(72) Inventors: **Yuh-shen SONG**, Porter Ranch, CA (US); **Catherine LEW**, Porter Ranch, CA (US); **Alexander SONG**, Porter Ranch, CA (US); **Victoria SONG**, Porter Ranch, CA (US)

(52) **U.S. Cl.**
CPC **G06Q 50/265** (2013.01); **G06Q 10/10** (2013.01); **G06F 16/285** (2019.01); **G06Q 30/0185** (2013.01); **G06F 16/22** (2019.01); **G06Q 20/4016** (2013.01)

(21) Appl. No.: **16/742,780**

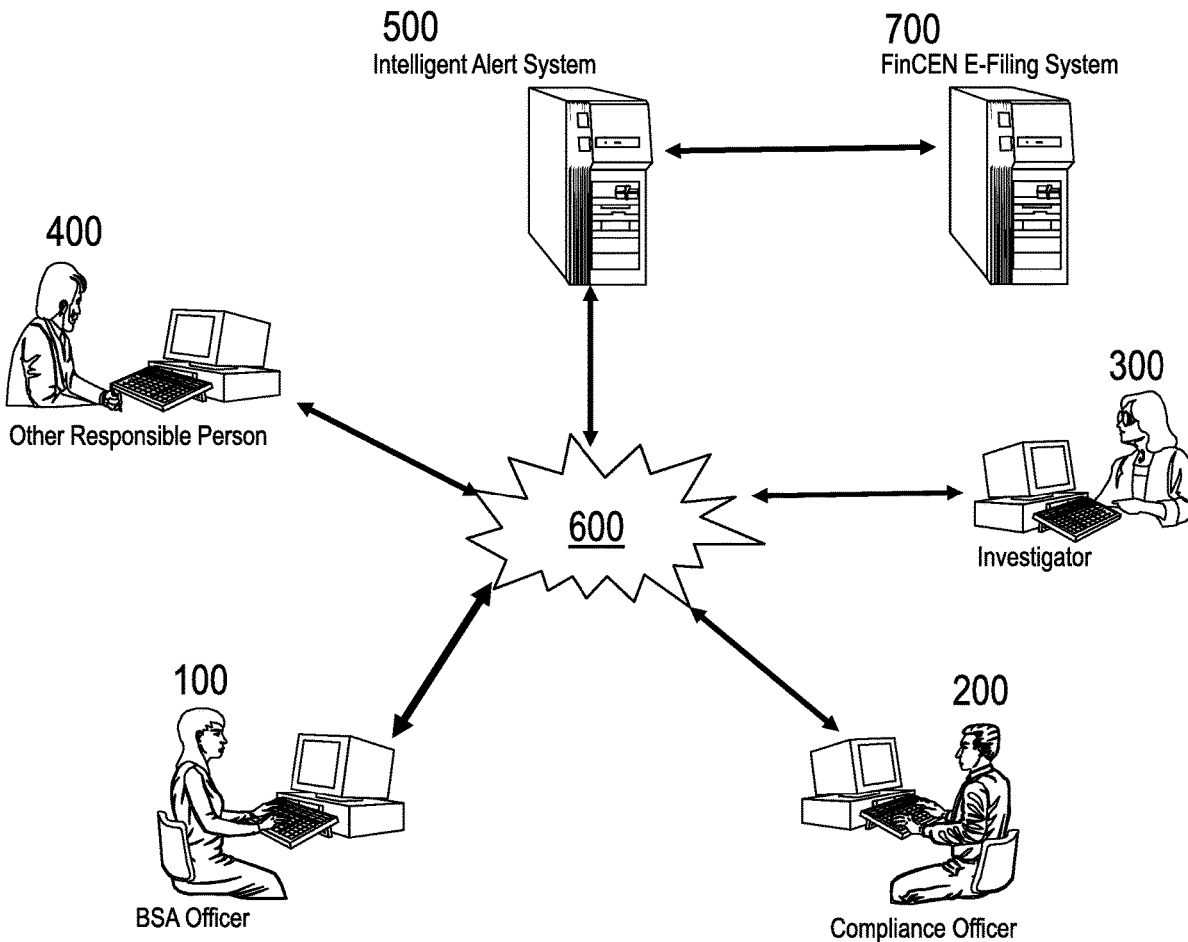
(57) **ABSTRACT**

(22) Filed: **Jan. 14, 2020**

An intelligent alert system triggers potential cases based on a set of scenarios and sends the potential cases and the reports with default narratives to an investigator. Through the changes made by the investigator on the default narratives, the intelligent alert system learns a writing style of the investigator. Eventually, the intelligent alert system can automatically submit a report with a narrative that matches the writing style of the investigator.

Related U.S. Application Data

(60) Provisional application No. 62/805,085, filed on Feb. 13, 2019.






FIG. 1

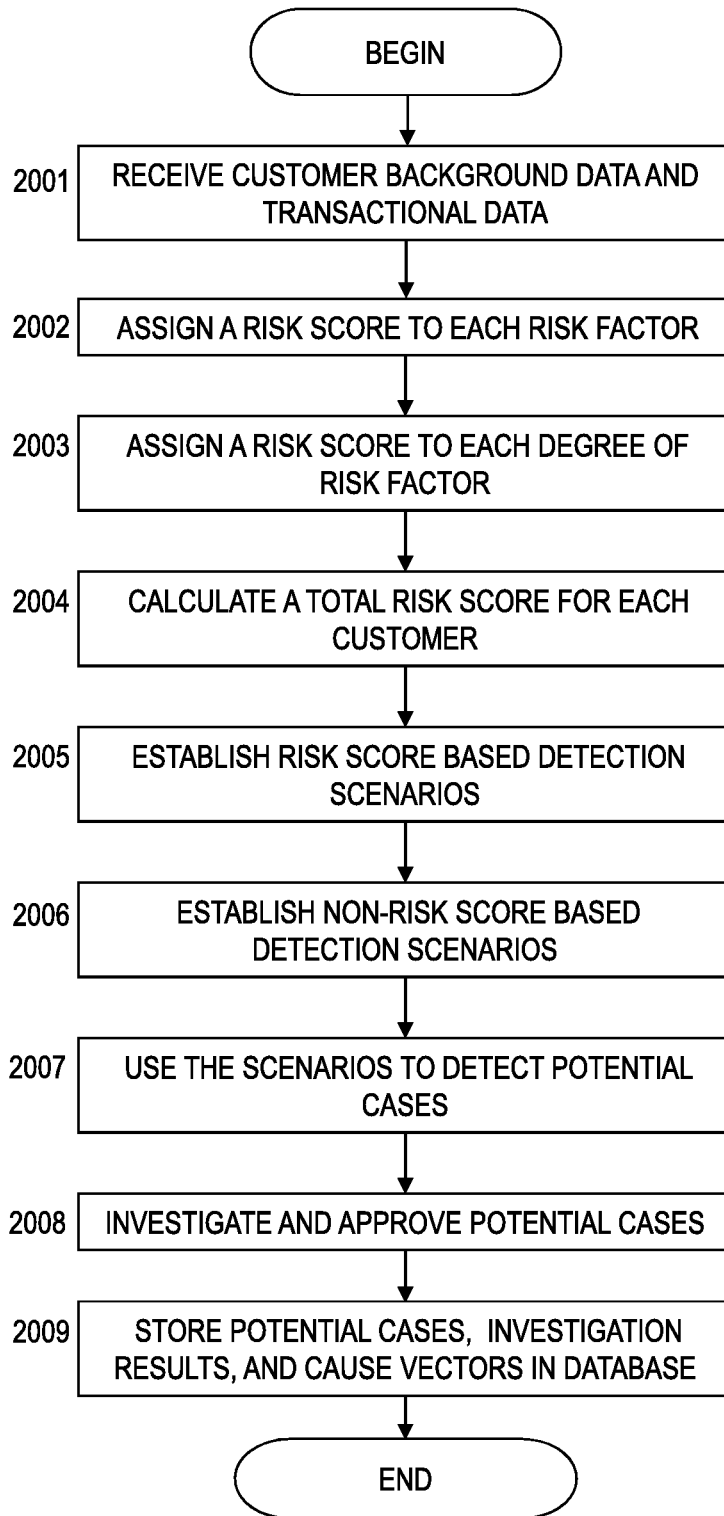


FIG. 2

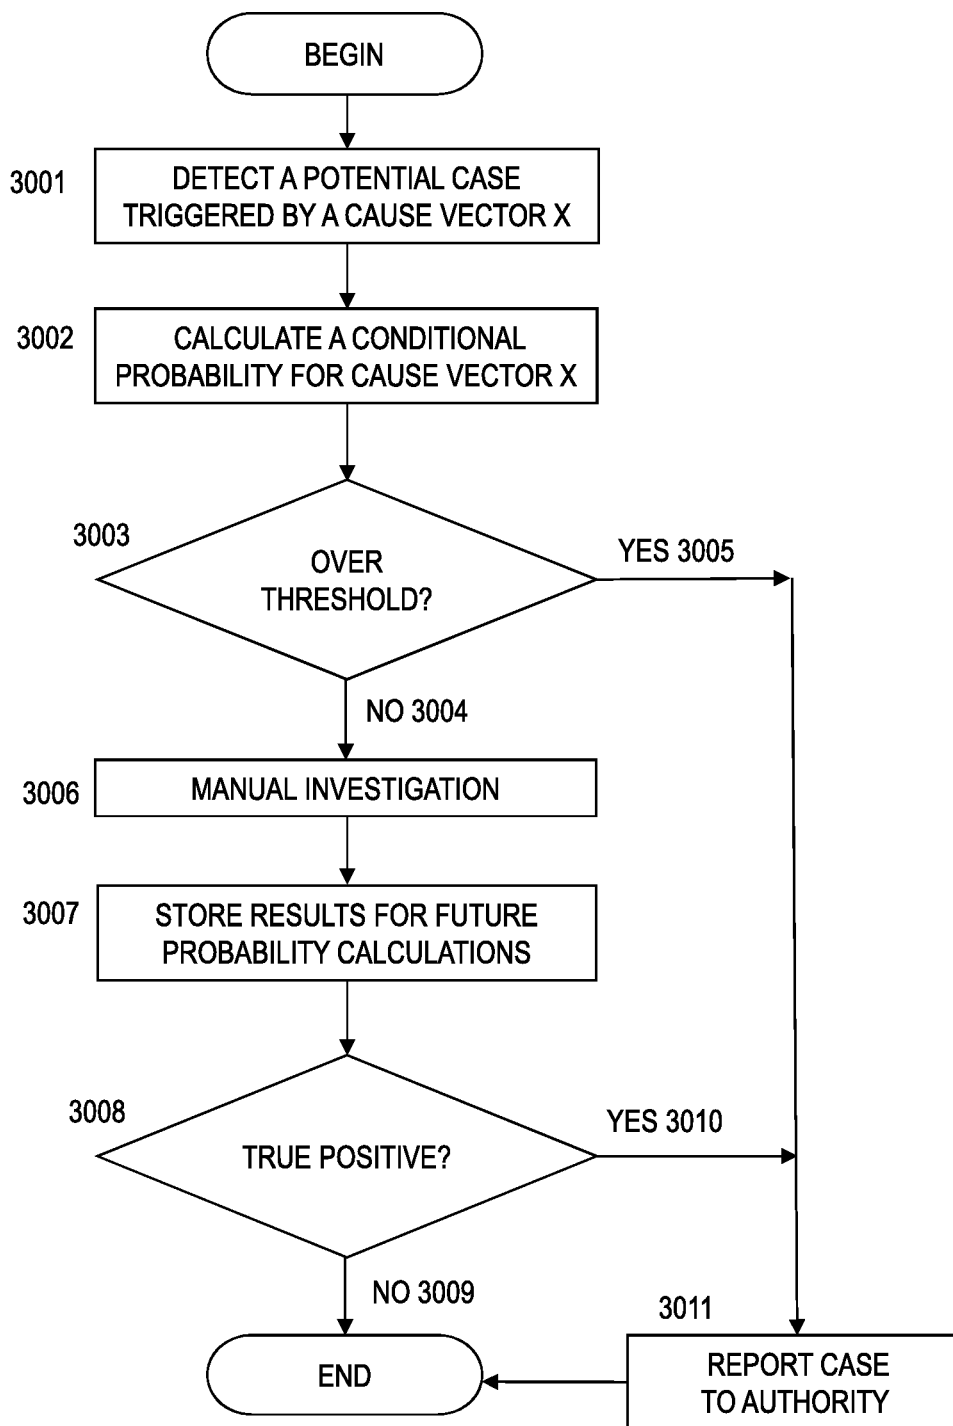


FIG. 3

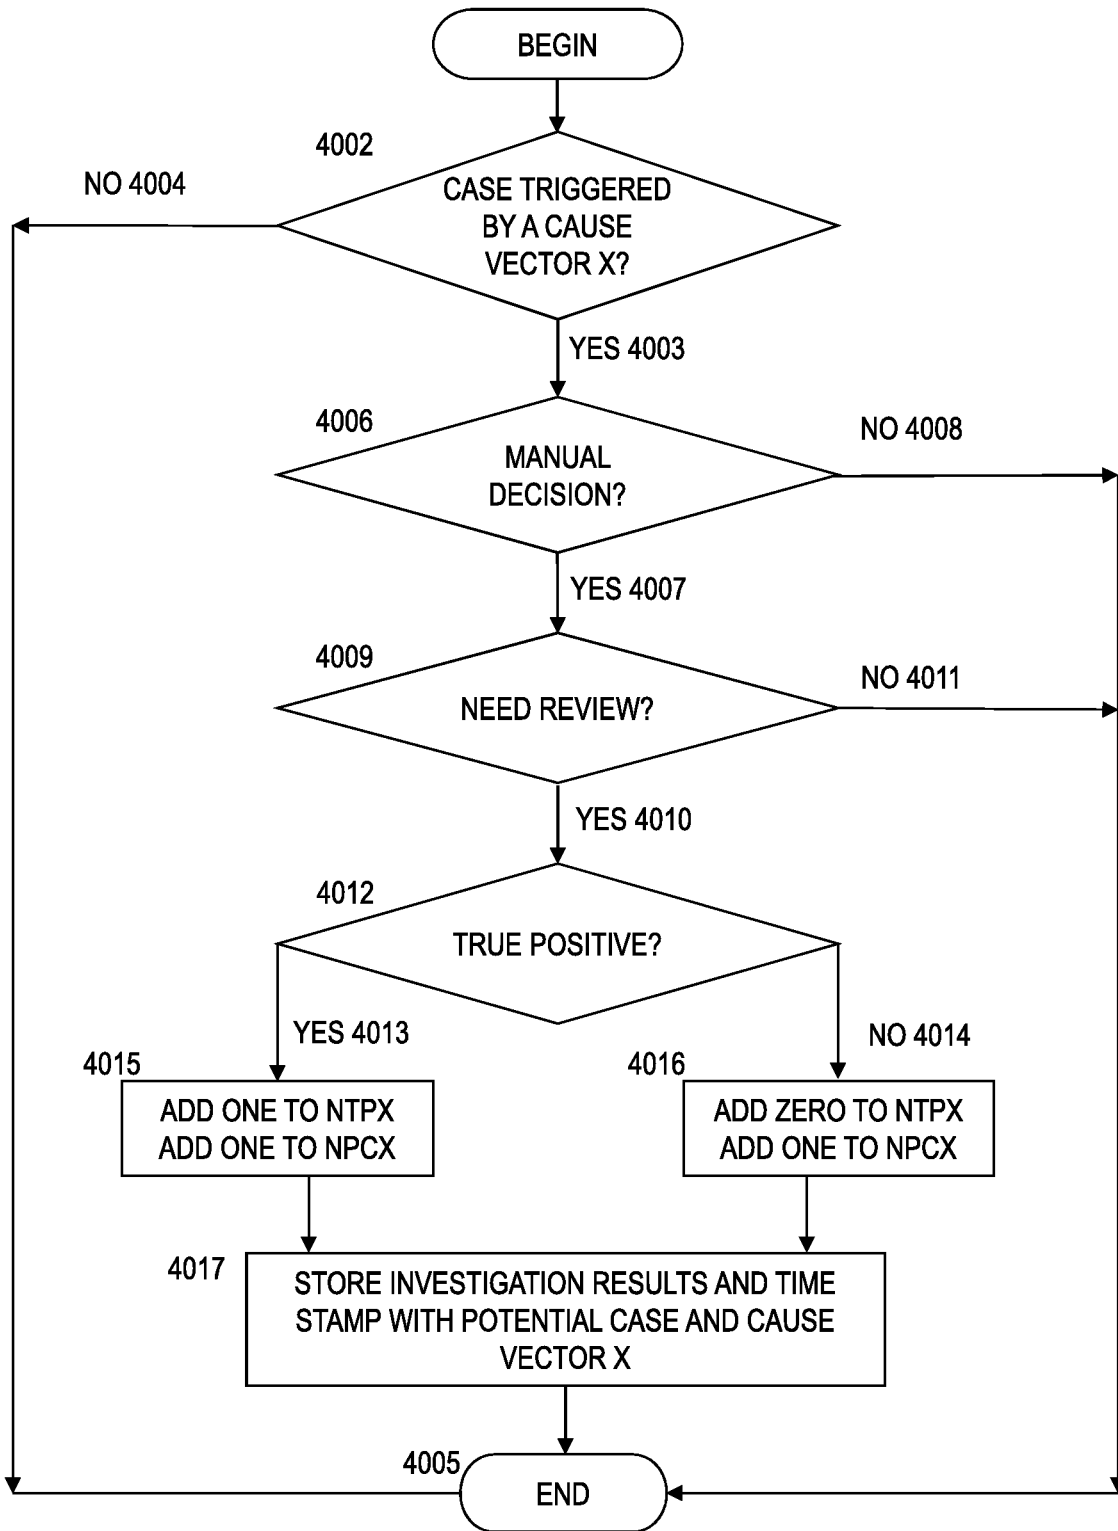


FIG. 4

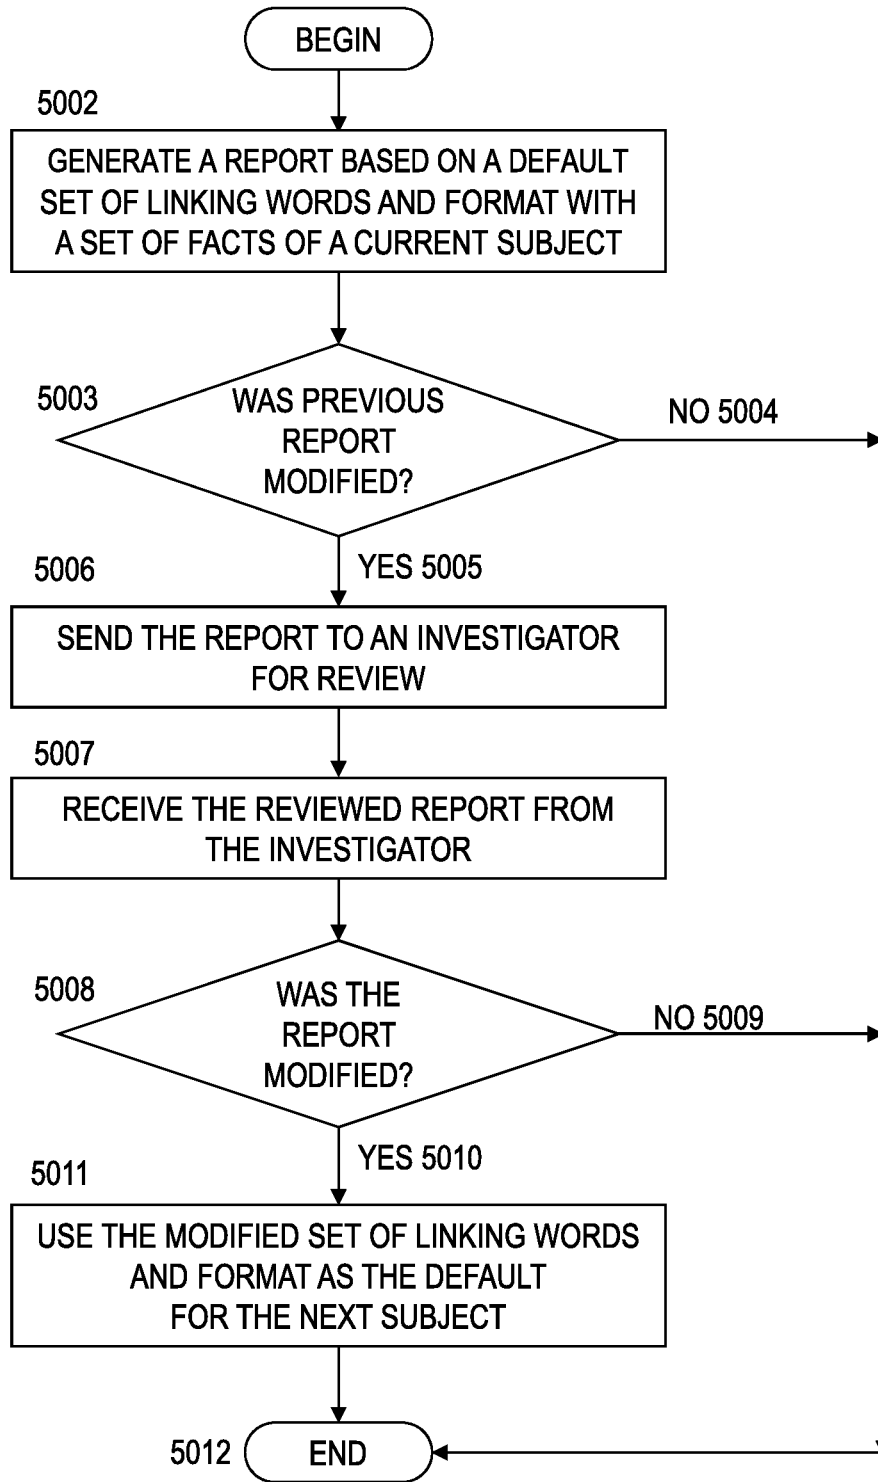


FIG. 5

INTELLIGENT REPORT WRITER

CROSS-REFERENCE TO RELATED APPLICATION

[0001] The present application claims the benefit of U.S. Provisional Patent Application No. 62/805,085, filed on Feb. 13, 2019, and entitled "INTELLIGENT ALERT SYSTEM," the disclosure of which is expressly incorporated by reference herein in its entirety.

TECHNICAL FIELD

[0002] The present disclosure relates generally to an intelligent alert system. More specifically, the present disclosure relates to a system and method for improving alert management.

BACKGROUND

[0003] The amount of data available for public consumption is increasing at an exponential rate. Data may be used to discover hidden opportunities or uncover bad incidents. Conventional information management systems may use manual searches, report based search systems, and/or alert based search systems. These conventional search systems may be used to detect and report suspicious activity.

[0004] The Bank Secrecy Act in the USA was first established in 1970. Under the Bank Secrecy Act, financial institutions must report suspicious activities to the government. Historically, financial institutions train frontline personnel (e.g., bank tellers) to observe and identify suspicious activities. Most financial institutions, however, could not effectively comply with the Bank Secrecy Act. After the 9/11 tragedy, U.S. lawmakers believed that effective compliance with the Bank Secrecy Act by financial institutions could have prevented the 9/11 tragedy.

[0005] To further enforce the Bank Secrecy Act, the U.S. Congress passed the USA PATRIOT Act, which enacted severe civil and/or criminal penalties for violations of the Bank Secrecy Act. Furthermore, the U.S. government agencies, such as Financial Crimes Enforcement Network (FinCEN), Office of Comptroller of Currency (OCC), Federal Reserve Bank (FRB), Federal Deposit Insurance Company (FDIC), National Credit Unions Administration (NCUA), State Banking Departments, Department of Financial Institutions, etc., strictly require financial institutions to comply with the Bank Secrecy Act, especially in their obligations to file Suspicious Activities Reports (SARs) to FinCEN.

[0006] Suspicious activities cover a very broad scope. For example, money laundering, terrorist financing, fraud, embezzlement, identity theft, computer intrusion, self-dealing, bribery, false statement, counterfeit instruments, mysterious disappearance, etc., are all classified as suspicious activities.

[0007] Nevertheless, many financial institutions have failed to detect and report suspicious activities. In fact, many financial institutions use products that are effective for preventing fraud, but ineffective for preventing money laundering or other financial crimes. In general, fraud can be detected based on a change of behavior because a fraudster that has stolen a victim's identity (or financial instrument) behaves differently from the victim. A computer system can detect a fraud case if an account's activities are different from expected activities as derived from historical activities.

[0008] For example, U.S. application (Publication No. 2003/0177087) specifies that a high risk variable can include a change in an account's usual behavior indicated, for example, when a transaction falls outside its profile. According to this publication, Beta, Delta, and Theta models are used to detect transactions that fall outside the profile of a customer.

[0009] However, money laundering and some other financial crimes can be committed without any change in behavior. As a result, the traditional approach of detecting fraud based on a change in behavior cannot detect some basic money laundering activities or other financial crimes. In the money laundering arena, a higher-risk customer may not be suspicious. For example, money services businesses (MSBs), pawnshops, ATM vendors, flight attendants, etc., are typically classified as higher-risk customers by banks in their anti-money laundering programs. Nevertheless, it does not mean that these higher-risk customers conduct money laundering activities. Although high risks are associated with these customers, there may be nothing wrong with these customers.

[0010] Some businesses are very difficult to monitor. For example, an MSB deals with a large number of transactions every day and a single money laundering transaction, mixed with a large number of transactions, may not be detected by the traditional approach.

[0011] The challenges noted for complying with the USA PATRIOT Act and the Bank Secrecy Act (BSA) are just some examples to illustrate the importance of identifying suspicious activities. Identifying suspicious activities can also be used to comply with other laws, such as the Fair and Accurate Credit Transactions Act (FACT Act), the Unlawful Internet Gambling Enforcement Act (UIGEA), the Elder Abuse Reporting Act, (EARA), the Sarbanes-Oxley Act (SOX), the regulations set by the Office of Foreign Assets Control (OFAC), and other laws and regulations.

[0012] Regulatory compliance is traditionally implemented through policies and procedures that require human workers to take some specific actions in response to certain conditions. For example, banks train their tellers in the branches to observe and report anything they see as suspicious to comply with the Bank Secrecy Act.

[0013] This traditional approach is no longer effective in the modern age because customers no longer need to appear in a branch of the bank. For example, customers can conduct remote electronic transactions (e.g., via the Internet) and there are many financial instruments available to customers (e.g., checks, credit cards, debit cards, etc.) Furthermore, perpetrators are sophisticated and know how to avoid attracting attention from tellers. As a result, depending on tellers to detect suspicious activities for compliance with the Bank Secrecy Act is insufficient.

[0014] Moreover, the cost of this human-based approach is very expensive. Intensive training must be conducted periodically to ensure that human workers truly know how to respond to each different situation in compliance with different laws and regulations. Human workers, however, are prone to mistakes. In fact, due to human oversight, many financial institutions have received severe penalties from government agencies for failure to comply with different laws and regulations.

[0015] It is desirable to improve search systems to improve the detection of different types of suspicious activity and help businesses comply with different types of laws

and regulations. The methods, functions, embodiments, computer systems, networks, software, hardware, mechanism, and other components used to detect suspicious activity may also be used for other applications, or other organizations for purposes other than detecting suspicious activity.

SUMMARY

[0016] This disclosure includes a number of embodiments that can be combined together to form a variety of methods. A method detects money laundering activity. The method includes detecting, by a first computer system, a first potential case for money laundering when flagged scenarios in a cause vector of the first potential case satisfy detection criteria. The method also includes comparing to a threshold, by the first computer system, a first ratio of a first value of the cause vector to a second value of the cause vector. The method also includes transmitting, from the first computer system to a second computer system, the first potential case for investigation when the first ratio is less than the threshold. The method further includes adjusting, by the first computer system, the first value when a result of the investigation indicates the first potential case is a true positive. The method also includes adjusting, by the first computer system, the second value based on the cause vector satisfying the detection criteria. The method still further includes transmitting, from the first computer system to a third computer system, a first report associated with the first potential case when the first potential case is the true positive.

[0017] Another method detects money laundering activity. The method includes detecting, by a first computer system, a potential case for money laundering when flagged scenarios in a cause vector of the potential case satisfy detection criteria. The method also includes calculating, by the first computer system, a conditional probability value for the potential case based on the cause vector. The method further includes comparing, by the first computer system, the conditional probability value with a threshold. The method also includes transmitting, from the first computer system to a second computer system, a report associated with the potential case when the conditional probability value is greater than the threshold.

[0018] Yet another method detects money laundering activity. This method includes detecting, by a first computer system, a potential case for money laundering when flagged scenarios in a first cause vector of the potential case satisfy detection criteria. The method also includes generating, by the first computer system, a combined cause vector by combining the first cause vector with a second cause vector of a prior potential case. The method further includes calculating, by the first computer system, a conditional probability value for a case triggered by the combined cause vector. The method also includes comparing, by the first computer system, the conditional probability value with a threshold. The method also includes transmitting, from the first computer system to a second computer system, a report associated with the potential case and the prior potential case when the conditional probability value is greater than the threshold.

[0019] Still another method detects money laundering activity. The method includes detecting, by a first computer system, a potential case for money laundering when flagged scenarios in a cause vector of the potential case satisfy

detection criteria. The method also includes calculating, by the first computer system, a conditional probability value for a case triggered by a child vector of the cause vector. The method further describes comparing, by the first computer system, the conditional probability value with a threshold. The method also describes transmitting, from the first computer system to a second computer system, a report associated with the potential case when the conditional probability value is greater than the threshold.

[0020] Another method detects money laundering activity. The method includes detecting, by a first computer system, a potential case for money laundering when flagged scenarios in a first cause vector of the potential case satisfy detection criteria. The method also includes generating, by the first computer system, a combined cause vector by combining the first cause vector with a second cause vector of a prior potential case. The method also includes calculating, by the first computer system, a conditional probability value for a child vector of the combined cause vector. The method further describes comparing, by the first computer system, the conditional probability value with a threshold. The method also describes transmitting, from the first computer system to a second computer system, a report associated with the potential case and the prior potential case when the conditional probability value is greater than the threshold.

[0021] A computer-implemented method generates a report. The method includes storing, in a database of a first computer system, a first fact associated with a first subject, a second fact associated with a second subject, and a third fact associated with a third subject. The first fact, the second fact, and the third fact have a same field name in the database. The method also includes receiving, at the first computer system from a second computer system, a first report of the first subject. The first report comprises the first fact and a first set of linking words generated by a human writer. The method further includes transmitting, from the first computer system to a third computer system, the first report of the first subject. The method transmits, from the first computer system to the second computer system, the second fact and the first set of linking words. The method further includes receiving, at the first computer system from the second computer system, a second report of the second subject. The second report comprises the second fact and a second set of linking words generated by the human writer. The method transmits, from the first computer system to the third computer system, the second report of the second subject. The method also transmits, from the first computer system to the third computer system, a third report of the third subject when the first set of linking words corresponds to the second set of linking words. The third report comprises the third fact and the second set of linking words.

[0022] The above methods are just some examples. Many other methods can be formed by combining and rearranging the embodiments of this disclosure.

[0023] This has outlined, rather broadly, the features and technical advantages of the present disclosure in order that the detailed description that follows may be better understood. Additional features and advantages of the disclosure will be described below. It should be appreciated by those skilled in the art that this disclosure may be readily utilized as a basis for modifying or designing other structures for carrying out the same purposes of the present disclosure. It should also be realized by those skilled in the art that such

equivalent constructions do not depart from the teachings of the disclosure as set forth in the appended claims. The novel features, which are believed to be characteristic of the disclosure, both as to its organization and method of operation, together with further objects and advantages, will be better understood from the following description when considered in connection with the accompanying figures. It is to be expressly understood, however, that each of the figures is provided for the purpose of illustration and description only and is not intended as a definition of the limits of the present disclosure.

BRIEF DESCRIPTION OF THE FIGURES

[0024] The features, nature, and advantages of the present disclosure will become more apparent from the detailed description set forth below when taken in conjunction with the drawings.

[0025] FIG. 1 illustrates a system and network diagram of an intelligent alert system according to aspects of the present disclosure.

[0026] FIGS. 2, 3, 4, and 5 are flowcharts for an intelligent alert system according to aspects of the present disclosure.

[0027] The detailed description set forth below, in connection with the appended drawings, is intended as a description of various configurations and is not intended to represent the only configurations in which the concepts described herein may be practiced. The detailed description includes specific details for the purpose of providing a thorough understanding of the various concepts. However, it will be apparent to those skilled in the art that these concepts may be practiced without these specific details. In some instances, well-known structures and components are shown in block diagram form in order to avoid obscuring such concepts. As described herein, the use of the term “or” may mean either “inclusive OR” or “exclusive OR,” depending on the situation of the application based on the convention.

DETAILED DESCRIPTION

[0028] Aspects of the present disclosure are directed to an intelligent alert system. In one configuration, a computer system monitors transactions and/or activity to generate alerts. The computer system may learn from humans and may become smarter to automatically accept potential cases as true positives and/or reject potential cases as false positives. As a result, this computer system can help a financial institution to reduce human resources while still complying with laws and regulations, such as the Bank Secrecy Act.

[0029] Depending on the specifics of a law or regulation, the computer system may use different functions or methods for monitoring different types of activity. The present disclosure provides various details for monitoring transactions and activity to reduce human resources while still complying with different requirements, laws, and regulations. The computer system of the present disclosure may also be used for other applications, or other purposes. The computer system of the present disclosure may reduce or eliminate human effort and/or mistakes, reduce resources, save money, and improve results.

[0030] In conventional information management systems, individuals manually search the Internet and/or databases for data. The manual searches are time consuming. To improve manual searches, decision makers often hire additional searchers to assist in performing manual searches.

[0031] Additionally, in some conventional information management systems, various reports with graphics are generated to summarize or compare data. Reading a report may be an improvement to manual searching. Still, reading the report may be time consuming when a large amount of data is involved. Additionally, it is impractical for human eyes to identify events amongst various numerical values presented in the report. A human may not identify different issues when tasked with reading a large number of reports.

[0032] To improve searches, some conventional systems generate an alert when a condition is satisfied. The alert system may mitigate the need for a human to read a report. That is, an alert may notify a user of a specific matter and provide data related to the matter. Compared with the report-based approaches, alert systems reduce an amount of time and manpower.

[0033] It is desirable to improve alert systems to increase efficiency and reduce the need for human oversight. Aspects of the present disclosure are directed to an information management system that processes alerts via a computer system. The information management system may be referred to as an intelligent alert system.

[0034] An alert system may be used by various types of organizations. For example, financial institutions are required, by law, to report suspicious activity to the government. Thus, financial institutions may use an alert system to determine when suspicious activity has been detected. As another example, loan companies may use an alert system to generate an alert when a borrower has a potential to default on a loan. In yet another example, social media companies may use an alert system to generate an alert when a cross-selling target is identified. As another example, defense contractors may use an alert system to identify a breach in security policy. In another example, police departments may use an alert system to generate an alert before a crime is committed. Alert systems are not limited to organizations and may also be used by families or individuals. For example, an individual may receive an alert when an investment opportunity materializes and/or before a stock market crash.

[0035] As discussed above, alert systems have various uses. Aspects of the present disclosure are not limited to the uses discussed above. The methods, functions, embodiments, computer systems, networks, software, hardware, firmware, mechanisms, and other components of the present disclosure may be used by other types of individuals and organizations for other purposes. For clarity, the present disclosure will discuss an example use of an alert system at a financial institution to detect suspicious activity.

[0036] The U.S. government strictly enforces a business's compliance with the USA PATRIOT Act, the Bank Secrecy Act (BSA), the Fair and Accurate Credit Transactions Act (FACT Act), the Unlawful Internet Gambling Enforcement Act (UIGEA), the Elder Abuse Reporting Act (EARA), the Sarbanes-Oxley Act (SOX), the regulations set by the Office of Foreign Assets Control (OFAC), and other related laws and regulations. The business may include, for example, financial institutions, such as banks, credit unions, mortgage companies, money services businesses, stockbrokers, and insurance companies. Billions of dollars in Civil Monetary Penalties (CMPs) have been levied by the U.S. government to financial institutions for violating these laws and regulations. Criminal penalties have also been issued to some individuals that work for the financial institutions.

[0037] A financial institution is just one type of business. Financial institutions are not the only organizations that need to comply with these laws and regulations. Many other types of businesses need to comply with these laws and regulations. The present disclosure applies to all businesses, such as businesses that are obligated to comply with laws and regulations.

[0038] The Financial Crimes Enforcement Network (FinCEN) and the Office of Foreign Assets Control (OFAC) are USA organizations. The laws and regulations in the USA are used as examples in the present disclosure. Many other countries have similar organizations that perform similar tasks. As such, similar laws and regulations exist in many other countries. The present disclosure is also applicable in those countries to help businesses comply with their respective laws and regulations. Aspects of the present disclosure may also be used by businesses, individuals, or organizations that do not need to comply with a law or regulation.

[0039] Often, it may be difficult to determine whether a person or a group of persons have conducted an illegal activity. According to the Bank Secrecy Act in the USA, when a business files a Suspicious Activity Report (SAR) with FinCEN, the business has no obligation to prove whether the reported case is an illegal activity. In fact, a “safe harbor” rule encourages businesses to report more suspicious activities without concern for the ramifications of being accused of reporting legitimate activities as wrongful by mistake. Under this “safe harbor” rule, no person (or organization) can file a lawsuit against an entity because the entity filed a Suspicious Activity Report (SAR) to FinCEN about this person (or organization). A SAR is used by the government to collect information, and a business is only expected to provide information and opinions in a SAR. Government agencies perform their own investigations to determine whether an activity reported in a SAR is truly illegal.

[0040] In general, the decision-making process regarding whether to report a suspicious activity, which is not fraud, is different from the decision-making process regarding whether to report a fraud case. For a fraud case, an entity, such as a business or a consumer, may lose money. Therefore, fraud is easier to detect in comparison to other crimes. As such, it is easier to decide whether to report a fraud case. Preventing fraud is also easier in comparison to preventing other crimes. As an example, if a computer system detects a high fraud risk associated with a transaction, the computer system can block the transaction and let an investigator investigate the transaction to determine whether it is truly a case of fraud.

[0041] In one aspect of the present disclosure, for fraud detection, a computer system calculates a risk score associated with a transaction based on different factors associated with the transaction. These factors may include historical activities of the account, deviations from the expected activities, location, time, amount, frequency and nature of the transaction, relationships between multiple accounts, type, nature and structure of the account holder, etc.

[0042] In one aspect of the present disclosure, for fraud detection, a computer system blocks the transaction if the fraud risk score of the transaction is over a threshold. The threshold can be pre-determined based on the policy of the business.

[0043] In one aspect of the present disclosure, for fraud detection, a computer system creates a case based on the

detected high fraud risk transaction. The case and related information are presented to an investigator for further investigation.

[0044] In comparison to fraud, there may be no clear evidence for suspicious activity. For example, a customer may frequently deposit a large amount of cash. It is possible that this customer is engaged in money laundering by selling illegal goods and receiving cash as payment. It is also possible that this customer sells homemade products in farmers’ markets and only accepts cash as payments. Often, due diligence is needed to determine whether there is anything suspicious.

[0045] It is also possible that although a customer sells homemade products in farmers’ markets, the customer also sells illegal goods at other locations. Unless the bank is informed that the customer sells illegal goods, there is no evidence for the bank to prove that the customer sells illegal goods. If the customer really sells illegal goods and the bank does not report such suspicious activity to FinCEN, the bank may receive a severe penalty later for failing to report the case to FinCEN once the customer is caught by the government for selling illegal goods.

[0046] On the other hand, if the bank reports every case that has a slight chance of being suspicious, the bank may invite unnecessary attention from government agencies. The government agencies may spend many months inside the bank to investigate the bank’s operations and may severely affect the bank’s operations.

[0047] The decision for reporting a case may be a judgment call by the person who reviews the case. In addition, the decision-making process can be quite subjective. Furthermore, a business cannot block a transaction simply because it appears to be a suspicious money laundering activity. A consumer may sue the business that blocks the consumer’s transaction when the business cannot really prove that money laundering has occurred. In fact, many government agencies often advise businesses, which have reported suspicious activities such as money laundering or terrorist financing, to keep quiet and process the suspicious transactions as normal transactions so that the suspect will not be alerted and flee. This approach gives government agencies more time and opportunity to identify all related perpetrators.

[0048] According to the U.S. Bank Secrecy Act, a business that files a SAR has the obligation to keep the SAR confidential and cannot let the suspect (e.g., a person involved in the case) know anything about the SAR, including the existence of the SAR. SARs can only be reviewed by authorized government agencies.

[0049] Because handling a suspicious activity case is very different from handling a fraud case, as described above, many conventional approaches and concepts that are applicable for fraud detection and prevention are no longer useful for detecting and managing suspicious activity, such as money laundering, terrorist financing, elder abuse, online gambling, etc. In one aspect of the present disclosure, a computer system records the opinions of the person that decides not to report a detected suspicious activity case. Under such circumstances, the decision maker records a reason to justify their decision.

[0050] Unlike a fraud case, a suspicious activity case may not be clear to a person who reviews the case until additional evidence becomes available. Therefore, it is possible that a person may dismiss a detected suspicious activity case at the

beginning, but change his mind later when additional evidence becomes available. In one aspect of the present disclosure, a person who reviews a detected suspicious activity case may also need to review all the historical detected cases regarding the same suspect to determine whether any new evidence, when combined with the old evidence, possibly from any dismissed cases, makes the newly detected case more suspicious. As a result, even if a case was previously dismissed as a false detection, such a dismissed case may be reviewed later.

[0051] This case review practice of suspicious activity may be different from the case review practice of fraud, as fraud cases typically have a clear conclusion. If a customer is a fraudster, the customer's account is closed and the customer will be prevented from conducting future transactions/activities. If a customer is a victim of fraud, the detected fraud case has nothing to do with the customer and the evidence is not used against the customer in the future. Therefore, a fraud investigator usually only focuses on newly detected case. Conversely, a suspicious activity investigator may need to review a history of the detected cases and make a decision after intensive research and analyses. In one aspect of the present disclosure, the justification of the decision for not reporting a suspicious activity is stored in a database and is available for future reference.

[0052] In another aspect of the present disclosure, a computer system also records the identity of the person that decides not to report the detected case. The computer system may compare the decisions made by multiple persons for not reporting suspicious activities of the same suspect(s) to determine whether an investigator is attempting to hide a detected suspect or case.

[0053] For a large business, thousands of suspicious activities may be detected every month. A group of persons may be tasked with reviewing the detected cases to determine whether the business needs to file SARs on these cases. In one aspect of the present disclosure, a computer system automatically allocates detected cases to different persons based on the policies set by the business. The computer system may monitor and record the status of each detected case. If a case review is delayed by a particular person, the computer system will alert the business of such delay.

[0054] In yet another aspect of the present disclosure, a computer system monitors the workload of each person who reviews detected cases. If a person has reviewed an unusually large number of cases in comparison to other persons who also reviewed detected cases during the same period of time, this person himself may become suspicious or questionable.

[0055] On the other hand, if a person has reviewed a small number of cases in comparison with other persons who also reviewed cases during the same period of time, this person may also become suspicious or questionable. In either of the above two situations, a manager of the business may want to investigate the situation and reach his own conclusion and resolution.

[0056] In general, different detection functions are used to detect suspicious activities because suspicious activities may arise in many different types of activities. Because the detection of suspicious activities is not clear, some detected cases may not be truly suspicious after investigation. Under such circumstances, such detected cases are dismissed as false detections or false positives. A false detection or false

positive is generally referred to as the conclusion of an investigation of a case, but not the reason to justify why the case is dismissed.

[0057] For example, if a financial institution detects a case where several customers live at the same address and deposit a large amount of cash into the financial institution, this case may be about a possible drug dealer family, with many of the family members depositing their proceeds from selling drugs. After investigation, however, this case may actually be a group of students living together and depositing tips that they receive from working at a restaurant. The reason to justify the decision not to report this case should be "students living together are depositing tips they have received from part-time jobs." As such, due to the given reason, the conclusion of the detected case becomes a false detection or false positive.

[0058] In general, after the review of a detected case, the case can be classified as a false detection (or false positive) by the person who reviewed this case. In one aspect of the present disclosure, a computer system provides information and/or statistics for a user to analyze all detected cases that have been classified as false detections. From these false detections, the user may identify detection functions that have generated a number of false detections that is greater than a threshold. The user can further improve the identified detection functions to improve the detection of future suspicious activities.

[0059] The USA PATRIOT Act, Bank Secrecy Act (BSA), Anti-money laundering (AML), and Anti-Terrorist Financing (ATF) have been important compliance matters in the financial industry since 9/11. Many financial institutions have invested large amounts of capital in these compliance matters, but still miss true money laundering and terrorist financing cases.

[0060] The primary cause for these compliance problems is that many financial institutions do not even detect basic money laundering cases, and senior managers of financial institutions have difficulty understanding these problems. Many financial institutions utilize a fraud detection principle for detecting money laundering activities and some even mix fraud cases and money laundering cases.

[0061] In reality, however, money laundering is very different from fraud. A fraud detection product can easily compare an account holder's current activities with the account holder's historical activities and detect possible fraud if the current activities deviate from the expected activities derived from the historical activities. For example, if a fraudster steals a credit card from a victim, the fraudster will conduct purchase activities that are different from the victim's historical activities. It is just a matter of time before the credit card company will detect the fraudulent activities and disables the credit card. If a new account does not yet have sufficient historical records, a fraud detection product compares the account holder's current activities with what the account holder said during the account opening process.

[0062] Because the goal of a fraud detection product is to stop losses as soon as possible, financial institutions usually run the fraud detection or risk-scoring in real time, or at least once daily. In contrast, real-time risk scoring, real-time detection, daily risk scoring, and daily detection methods that are effective for fraud detection cannot detect many basic money laundering activities. In fact, as explained earlier, a higher-risk customer may not be a money laun-

derer. It is a waste of time to assume that higher-risk customers are conducting suspicious money laundering activities.

[0063] A financial institution typically has a Bank Secrecy Act Officer (BSA Officer) who is responsible for reporting suspicious money laundering or terrorist financing activities to FinCEN. The following case is an example of how BSA Officers inside financial institutions can waste a significant amount of time reviewing their real-time risk scoring or daily risk scoring results, while still missing true money laundering cases. This example consists of the following facts: (a) Client A sends less than \$3,000 to XYZ around the 5th day of each month; (b) Client B sends less than \$3,000 to XYZ around the 8th day of each month; (c) Client C sends less than \$3,000 to XYZ around the 12th day of each month; (d) Client D sends less than \$3,000 to XYZ around the 17th day of each month; (e) Client E sends less than \$3,000 to XYZ around the 24th day of each month; (f) Client F sends less than \$3,000 to XYZ around the 29th day of each month; (g) A, B, C, D, E and F are unrelated individuals; and (h) XYZ is a drug dealer in Los Angeles with no prior criminal record.

[0064] In the above example, if a BSA Officer compares a client's current activities with the client's historical activities to detect any change of behavior, the BSA Officer does not detect anything unusual because the clients have consistently conducted similar transactions each month. If the bank tellers ask the clients about the purpose of the fund transfers, the clients can easily lie. Because these clients conduct their transactions on different days throughout the month, a BSA Officer would not be able to detect any risk on any given day of the month.

[0065] Furthermore, these clients are not related, and therefore the BSA Officer would not see their aggregate activities. In addition, because each transaction only involves a small dollar amount occurring once a month and the recipient of the funds resides in a U.S. city with a large population and heavy commercial activities, none of these clients would be viewed as high risk or suspicious based on these transactions. As a result, a fraud detection product will miss these basic money laundering cases despite the fact that the BSA Officer is working diligently with the fraud detection product every day.

[0066] To detect these money laundering cases, in one configuration, a computer system collects transactional data from the financial institution and conducts data mining based on anti-money laundering and anti-terrorist financing scenarios across all the transactions of all clients for a specified period of time, such as 30 days or longer. A computer system may collect all funds transfer transactional details from different data sources, such as wire, ACH, card payments, mobile payments, etc., inside a financial institution. The computer system may then identify a common recipient of these funds transfer transactions.

[0067] When a common recipient is identified, the computer system can display all transactions sent to a common recipient to the BSA Officer. The BSA Officer reviews the identified transactions through the computer system. The BSA Officer also reviews all historical cases associated with the suspects of the newly detected case. If the BSA Officer (e.g., a responsible person) agrees that such transactions are suspicious activities because the common recipient receives too much money, the computer system assists the BSA Officer to file a SAR to FinCEN. If the BSA Officer decides

not to file a SAR, the BSA Officer enters a reason into the computer system to justify his decision for not reporting such detected activities.

[0068] There are several methods to report SAR cases to FinCEN. One approach is to send the SAR report in electronic format directly to the server located at FinCEN. Under such circumstances, a BSA Officer can instruct the computer system that has detected suspicious activities to file a SAR report. The computer system will prepare the SAR report based on the suspect and the transactions identified by the BSA Officer, then transmit the SAR report to the computer system at FinCEN.

[0069] As one can understand, data mining a huge amount of transactional data of all clients of a financial institution accumulated over a long period takes some time even for a very small financial institution. Because a financial institution will not directly lose any money in a money laundering case, according to the regulatory guidelines, a BSA Officer has up to 30 days to file a SAR. This example illustrates that it is a waste of time and resources to conduct real-time risk scoring or daily risk scoring that actually misses true money laundering activities.

[0070] BSA Officers voice a common frustration that they are wasting time on false positives every day, at the expense of detecting actual money laundering cases. This frustration is the result of the prevailing misconception that money laundering and fraud are often crimes committed by the same offender and should be detected together based on detected change in behavior. After purchasing fraud detection products, some financial institutions try to detect both money laundering cases and fraud cases together. This has resulted in a huge amount of time, money, and resources being wasted. This misconception can be corrected through a proper understanding of the sophisticated facets of transactional risks.

[0071] Transactional risks are defined as risks directly associated with the transactions. For example, money laundering risk and fraud risk are directly associated with the transactions. Nevertheless, these risks possess very different characteristics. Customers who conduct money laundering through financial institutions intend to use the financial institutions as vehicles to achieve their goals. These money launderers usually pretend to be good customers because they need the financial institutions' assistance to accomplish their schemes. They do not mind paying extra fees or losing interest on their money, and thus from the financial institutions' perspective, these money launderers are desirable customers. This is one of the key reasons why financial institutions need to conduct data mining on all transactions to detect money laundering activities that are hidden behind the scene.

[0072] In comparison, fraud risks manifest themselves differently. Fraud committed by customers is generally classified into two categories: (1) third-party fraud; and (2) counter-party fraud. Third-party fraud is defined as fraud committed by a third party that is not the financial institution and is not the customer. For example, both the financial institution (e.g., primary party) and the customer (e.g., counter party) may become victims when a fraudster (e.g., third party) steals a checkbook from the customer. Under such circumstances, the transactions conducted by the third-party fraudster have nothing to do with the customer. It is therefore a waste of time, money, and resources when BSA Officers are misled by an ineffective fraud detection product

to assume that a customer has conducted money laundering (e.g., when there is a change of behavior) simply because the customer is a victim of fraud committed by a third party.

[0073] Counter-party fraud is defined as fraud committed by a customer (e.g., counter party) who cheats the financial institution (e.g., primary party). Once the customer has successfully cheated the financial institution, the customer quickly disappears and does not conduct money laundering through the financial institution. A fraudster may use financial institution A to launder the money, which the fraudster has stolen from financial institution B. To financial institution B, this is a fraud case. To financial institution A, this is a money laundering case. Neither of financial institution A nor financial institution B, however, sees both the fraud case and the money laundering case occurring to this same customer. Clearly, a system that intends to detect fraud cases every day systematically creates many false positives for money laundering and actually misses the real money laundering cases. Using such an approach increases the workload of the BSA Officers and exposes the financial institution to unnecessary regulatory risk.

[0074] There are other risks under the category of third-party fraud. For example, counterfeit checks, credit card fraud, debit card fraud, ATM fraud, online fraud, etc., are typical risks under the category of third-party fraud. Similarly, there are many different risks under the category of counter-party fraud, such as check kiting, deposit fraud, loan fraud, etc. Therefore, a good transactional risk management system uses multiple detection algorithms that intelligently take into account each unique characteristic of the various types of fraud to successfully detect fraud.

[0075] Furthermore, as explained earlier, multiple customers may launder money or finance terrorists together by conducting one small transaction for each person on different days, and daily monitoring misses such cases. This leads to the logical conclusion that a system using a single method to detect behavior changes wastes resources and misses true money laundering and terrorist financing cases. In one aspect of the present disclosure, money laundering and terrorist financing activities are detected by a different detection method, which conducts data mining on all transactions of the entire financial institution accumulated over a period of time based on a user-defined scenario.

[0076] In one aspect of the disclosure, a computer system uses multiple detection methods for monitoring transactions and integrates the detection results into a centralized case management platform. This approach consolidates and streamlines anti-money laundering, anti-fraud, and anti-financial crimes to improve detection while maintaining a holistic, accurate picture at all times. As a result, a financial institution can improve compliance with the regulatory requirements, eliminate risks, avoid losses, improve productivity, reduce resources used for managing transactional risks, reduce costs associated with hardware, database and software, lower IT maintenance workload, and increase overall profitability.

[0077] In one aspect of the present disclosure, a computer system compares a transactional pattern of a customer (or a group of customers) with known money laundering transactional patterns to detect suspicious money laundering activities. If there is a match, a possible money laundering activity may have been detected.

[0078] For example, many perpetrators know that a bank has to file a Currency Transaction Report (CTR) to the U.S.

government if more than \$10,000 in cash is deposited into a bank account on the same day. To avoid the filing of CTRs, perpetrators often split one large cash deposit into multiple smaller cash deposits, each cash deposit occurring on a different day and each cash deposit being less than \$10,000. This transactional pattern is called “structuring,” a known money laundering transactional pattern, and a computer system can detect this type of transactional pattern. There are many other types of transactional patterns that are known as money laundering transactional patterns. A computer system can be designed to detect each of these known money laundering transactional patterns. As a result, even if there is no change of behavior, a money laundering activity can be detected based on the transactional pattern of the suspect or suspects.

[0079] In one aspect of the present disclosure, the BSA Officer (or responsible person) investigates the detected case to determine whether it is a true money laundering case. In one aspect of the present disclosure, the BSA Officer also reviews all historical cases associated with the suspect(s) of the currently detected case. In one aspect of the present disclosure, if the BSA Officer agrees that such transactions are suspicious activities, the computer system assists the BSA Officer to file a SAR to FinCEN. In another aspect of the present disclosure, if the BSA Officer decides not to file a SAR, the BSA Officer enters a reason into the computer system to justify his decision for not reporting such detected activities.

[0080] In another aspect of the present disclosure, a group of customers with one or more common risk factors (or characteristics), such as a type of business, business model, organization structure, size, location, product, service, career type, position, etc., are compared together to detect suspicious money laundering activities. If a transactional activity (e.g., transactional pattern, transactional volume, transaction frequency, transaction trend, number of transactions, amount of transactions, transaction derivative, etc.) of one customer is different from the transactional activities of other customers, this customer may have conducted a suspicious money laundering activity. In one aspect of the present disclosure, statistical figures such as mean, variance, standard deviation, etc. of the group of customers are used to facilitate such comparison. Similarly, if one customer behaves differently from other customers who have the same set of risk factors (or characteristics), this customer may have conducted a suspicious money laundering activity. As a result, even if there is no change in behavior in any account, a suspicious money laundering activity can be detected.

[0081] Sometimes, it is not easy to compare a group of customers together. For example, an MSB with 100 branches may have many more cash activities than another MSB with only two branches. In one aspect of the present disclosure, to achieve a more effective comparison, it is useful to compare some derivatives (e.g., ratio of several numbers) instead of the original raw data. For example, a ratio can be “total cash withdrawal from a bank divided by total number of checks deposited into a bank.” In this example, the number of checks deposited can be used to measure the size of the MSB’s check-cashing operations. Therefore, the ratio “total cash withdrawal divided by total number of checks deposited” basically scales the check cashing operations of the 100-branch MSB and the check cashing operations of the two-branch MSB to about the

same level, based on check-cashing activities so that they can be compared on a more level ground.

[0082] Many other derivatives can be used to achieve a better comparison. In general, a derivative for a more effective comparison may include “a first variable of interest divided by a second variable which measures the size of the business (or operations).” For example, “total ACH outgoing transactional amount divided by total number of checks deposited,” “total wire outgoing transactional amount divided by total number of checks deposited,” “total number of prepaid cards issued divided by total number of checks deposited,” “total ACH outgoing transactional amount divided by total number of branches,” “total wire outgoing transactional amount divided by total number of branches,” “total number of prepaid cards issued divided by total number of branches,” “total ACH outgoing transactional amount divided by total number of prepaid cards issued,” “total wire outgoing transactional amount divided by total number of prepaid cards issued,” etc., are just some examples of the possible derivatives that can be used. In one aspect of the present disclosure, in addition to the ratios above, other forms of mathematical transformation create a derivative.

[0083] In one aspect of the present disclosure, a computer system compares a derivative of a particular customer with the derivatives of a group of customers who have one or more common risk factors (or characteristics) with the particular customer (e.g., same type of business or profession). If the derivative of the particular customer noticeably deviates from the derivatives of the group of customers, the particular customer may have conducted a suspicious money laundering activity. In one aspect of the present disclosure, statistical analysis such as mean, variance, standard deviation, etc., of the group of customers facilitate such comparison.

[0084] In one aspect of the present disclosure, a computer system uses many different risk factors to determine the money laundering risk of each customer of a financial institution. For example, these risk factors may include an industry, category of the customer, business type of the customer, geographical area of the customer, country of the address of the customer, a nature of the customer’s business product types of the business, services types of the business, structure of the business, profession of the customer, nationality, historical records (including the compliance records such as the number of Currency Transactional Reports, the number of Suspicious Activity Reports, the matches with the OFAC list, the matches with the 314(a) list, the matches with the Political Exposed Persons list, special designations by the compliance program, etc.), type of the transaction conducted, balance of the accounts, funds inflow, funds outflow, transactional pattern, number of transactions, amount of transactions, transactional volume, transactional frequency, transactional derivative, location of the transaction, time of the transaction, country of the transaction, sender of a money transfer transaction, location of the sender, country of the sender, nature of the sender, recipient of a money transfer transaction, location of the recipient, country of the recipient, nature of the recipient, relationships, social status, political exposure, historical transactions, etc. In fact, thousands of risk factors may be considered to determine the money laundering risk of a customer. For the purpose of this

disclosure, a “risk factor” is also referred to as a “representative element of a risk dimension,” or simply a “risk dimension.”

[0085] According to aspects of the present disclosure, each attribute of a customer that can affect the risk of the customer is a risk factor. Additionally, each characteristic of a customer that can affect the risk of the customer may be a risk factor. Furthermore, each type of activity of a customer that can affect the risk of the customer is a risk factor. The risk factors may also be affected by other risks, such as a piece of information related to a customer, each type of transaction of a customer, and/or each transactional pattern of a customer. Each risk factor is assigned a risk value.

[0086] In one configuration, each degree of the same type of risk is a risk factor and is given a risk score. For example, a total cash transactional amount in a period of 30 days can be used to measure the degree of the risk associated with money laundering. For example, we can define a total cash transactional amount level (or total cash transactional amount degree) from \$0 to \$5,000 during a 30-day period to have a risk score of 10; from \$5,001 to \$50,000, a risk score of 50; from \$50,001 to \$250,000, a risk score of 100; from \$250,001 to \$1,000,000, a risk score of 200; from \$1,000,001 to \$10,000,000, a risk score of 500; and for \$10,000,000 and above, a risk score of 1,000. In this example, a person with a total cash transactional amount of \$60,000 during a period of 30 days is classified into the amount level of “between \$50,001 and \$250,000,” and has a risk score of 100.

[0087] “Cash transaction amount” is just used as an example. Other considerations, such as number of cash transactions, acceleration of cash transaction, etc., can also be used to measure the degrees of risks associated with money laundering. In addition to cash, other financial transactions such as checks, wires, ATM, ACH, virtual currency, virtual securities, virtual instruments, credit card, debit card, prepaid card, monetary instruments, transfer, etc., can also be used to measure the degrees of risks associated with money laundering. Those skilled in the art can easily figure out numerous risk factors based on the above example.

[0088] In one aspect of the present disclosure, a risk score based scenario is based on customer data. Each piece of information about the customer is a risk factor and is assigned a risk score. Additionally, or alternatively, a risk score based scenario is based on transactional data. Each amount level (or amount degree) of a type of transaction is a risk factor and is assigned a risk score.

[0089] In one aspect of the present disclosure, customer data is associated with one or more of an industry category of the customer, a business type of the customer, a geographical area of the customer, a country of an address of the customer, a nature of a business of the customer, a product type of the business, a services type of the business, a structure of the business, a profession of the customer, a nationality of the customer, a historical record, a type of the transaction conducted, a balance of an account, funds inflow, funds outflow, a transactional pattern, a number of transactions, an amount of transactions, a transactional volume, a transactional frequency, a transactional derivative, a location of the transaction, a time of the transaction, a country of the transaction, a sender of a money transfer transaction, a location of the sender, a country of the sender, a nature of the sender, a recipient of a money transfer transaction, a location of the recipient, a country of the recipient, a nature of the

recipient, a relationship, social status, political exposure, a historical transaction, a number of suspicious activity reports (SARs) filed for money laundering and terrorist financing cases, a category of a first financial institution, a business type of the first financial institution, geographical area of the first financial institution, country of a head office of the first financial institution, nature of the business of the first financial institution, age of a person, sex of the person, income level of the person, appearance of the person, judgment about the person, a personal condition of the person, a family condition of the person, a family member of the person, a family member's condition of the person, a friend of the person, a friend's condition of the person, a historical record of the person, an industry category of the person, a geographical area of the person, a country of an address of the person, a profession of the person, a job type of an employee, an education level of an employee, an income level of an employee, a length of employment at a current job, a performance review record, employment history, a duration of each employment in the employment history, a reason for termination of each employment in the employment history, an age of the employee, a sex of the employee, a personal condition of the employee, a family condition of the employee, a family member of the employee, a family member's condition of the employee, a friend's condition of the employee, a historical record of the employee, a type of work performed, a number of transactions performed, an amount of transactions performed, a largest amount of transaction, a number of transactions with a particular counter party, an amount of transactions with a particular counter party, a number of changes of a crucial record, a number of changes of a crucial record associated with a particular counter party, a geographical area of an employee's home, a geographical area of an employee's office, a country of the address of the employee, a due diligence result of the customer, a length of an account history, a number of name matches with gambling organizations in transactions, or a combination thereof.

[0090] In one aspect of the present disclosure, transactional data is associated with one or more of cash, check, wire transfer, ATM (Automated Teller Machine), ACH (Automated Clearing House), virtual currency, virtual securities, virtual instruments, credit card, debit card, prepaid card, electronic fund transfer, wires, monetary instruments, letters of credit, notes, securities, commercial papers, commodities, precious metal, account opening, account closure, an account application, deposit, withdrawal, cancellation, balance check, inquiry, credit, debit, or a combination thereof.

[0091] In one aspect of the present disclosure, each risk factor is given a risk score and a customer is given a total risk score, which is a summation of all risk scores of the risk factors associated with the customer. This process of producing a total risk score for each customer may be referred to as risk scoring. This total risk score is used to determine the level of risk associated with the customer. A summation is used as one example in this disclosure. In fact, many different types of mathematical transformations can also be used to achieve a similar effect.

[0092] In one aspect of the present disclosure, each risk factor is given a risk score and a customer is given a total risk score, which is a value derived from a mathematic transformation of all risk scores of the risk factors associated with the customer.

[0093] As explained earlier, unlike a fraud situation, a higher-risk client may not be a suspect for money laundering or terrorist financing. High risk may just be the nature of the client. For example, MSBs, pawnshops, car dealers, pilots, flight attendants, etc., are often classified as higher-risk customers for anti-money laundering and anti-terrorist financing purposes, however, this does not mean that these customers are conducting money laundering activities or terrorist financing.

[0094] Nevertheless, because a customer has a high risk score, the customer may be closely monitored and a different monitoring method may be applied. Therefore, in one aspect of the present disclosure, the total risk score of a customer is used to determine the monitoring method applied to monitor the customer. If a total risk score of a customer is higher, a closer monitoring method is applied to monitor the customer. If a total risk score of a customer is lower, a more relaxed monitoring method is applied to monitor the customer.

[0095] In other words, in one aspect of the present disclosure, a total risk score of a customer is not used to determine whether the customer is suspicious. Instead, a total risk score of a customer is used to select an algorithm or a set of algorithms for monitoring the customer.

[0096] Sometimes, a customer with a very high risk score may be suspicious. Therefore, in one aspect of the present disclosure, if the total risk score of a customer is higher than a predefined value, an alert about the customer will be triggered so that the investigator can investigate the potential case. The predefined value can be set by a software module, a person who designs the system, a person who tunes the system, a person who uses the system, or a combination thereof.

[0097] In one aspect of the present disclosure, a group of customers having the same risk factor are compared together. For example, we can compare all the customers who are flight attendants together. In one aspect of the present disclosure, if the total risk score of a particular flight attendant is much higher than a reference value derived from the total risk scores of all flight attendants, this particular flight attendant may have conducted some suspicious money laundering activities. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values.

[0098] Statistical approaches can also be applied to facilitate the detection of suspicious activities. For example, mean, variance and standard deviation can be derived from the total risk scores of all the customers who are flight attendants. In one aspect of the present disclosure, if the total risk score of a particular flight attendant is higher than the mean of total risk scores of all flight attendants by over 4 times the standard deviation, this particular flight attendant may have conducted suspicious activities.

[0099] The above reference "4 times" is just one example. The number "4" can be any number, such as 3.75, 4.21, 10, etc. In one aspect of the present disclosure, if the total risk score of a particular flight attendant is higher than the mean of total risk scores of all flight attendants by over x times of the standard deviation, this particular flight attendant may have conducted suspicious money laundering activities, where x is a number assigned by the BSA Officer (or a responsible person). This statistical approach can be applied whenever a group comparison is used.

[0100] A flight attendant is merely one example to illustrate this method to detect suspicious money laundering activities among a group of entities. In practice, many other risk factors can be used for similar purposes. Because there are tens of thousands of risk factors, in one aspect of the present disclosure, a computer system permits a user to select any risk factor to identify all customers who have the same risk factor. In one aspect of the present disclosure, if a particular customer has a total risk score which is much higher than a reference value derived from the total risk scores of other customers who have the same risk factor, the particular customer may have conducted suspicious money laundering activities. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values.

[0101] Instead of one risk factor, a group of risk factors can also be used. In fact, a group of risk factors may improve the accuracy of the detection results. For example, in addition to the risk factor of profession (e.g., flight attendant), the destination country of the flight on which the flight attendant works can be another useful risk factor to detect money laundering risk. For example, a flight attendant who works on a flight between New York and Chicago may have activities, which are different from the activities of another flight attendant who works on a flight between Miami and Mexico City. It may be more accurate to compare a subgroup of flight attendants who work on flights between Miami and Mexico City. In this example, two risk factors, profession and destination city of flight, are considered to improve the accuracy of detection.

[0102] In one aspect of the present disclosure, a set of risk factors is used to identify a group of entities. If a particular entity has a total risk score that is much higher than a reference value derived from the total risk scores of all entities having the same set of risk factors, the particular entity may have conducted suspicious money laundering activities. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values. To simplify the calculations, the standard group statistics such as mean, variance, standard deviation, etc., which can be easily calculated based on existing software developing tools, can be derived to facilitate such comparison among a group of entities. As a result, even if there is no behavior change in any account, a computer system can still detect suspicious money laundering activities based on the above approach.

[0103] Sometimes, it may be helpful to eliminate some entities from the group comparison process because such entities are very different from others. In one aspect of the present disclosure, a computer system permits users to select some entities that will not be included in a group comparison process.

[0104] Detecting a flight attendant as having a suspicious money laundering activity is just one example. A similar method can be applied to many other different situations. For example, it is usually very difficult for banks or credit unions to detect a Money Services Business (MSB) customer as having suspicious money laundering or terrorist financing activity because an MSB has many transactions every day and one money laundering transaction may be hidden among many other normal transactions.

[0105] In one aspect of the present disclosure, an additional risk factor (e.g., near Mexico border) is used to identify a group of MSBs (e.g., in addition to the first risk

factor—type of business) which have this same set of risk factors. If a particular MSB has a total risk score higher than a reference value derived from the total risk scores of all MSBs having the same set of risk factors, the particular MSB may possibly have conducted suspicious money laundering activities. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values. Similarly, the standard group statistics such as mean, variance, standard deviation, etc. can be derived to facilitate such comparison among a group of MSBs.

[0106] Sometimes, it is not easy to compare a group of MSBs because they may have different types of operations and different sizes. In one aspect of the present disclosure, a part-time MSB and a full-time MSB are given two different risk factors because they may have different natures of business. In another aspect of the present disclosure, each of the different types of MSB products and/or services is given a risk factor. For example, each of money transfer, check cashing, currency exchange, prepaid card management, etc. is given a risk factor, although all of them can be offered by the same MSB. In one aspect of the present disclosure, a set of risk factors that precisely defines the types of products and/or services is used to identify the risks.

[0107] In one aspect of the present disclosure, some risk factors are adjusted based on the size of the operations so that a group comparison becomes more effective. For example, an MSB with 50 branches may naturally have 5 times the total cash transactional amount of another MSB with 10 branches. Sometimes, to conduct a group comparison, risk factors affected by the size of the operations may be adjusted to account for the size of the operations. For example, for an MSB with 50 branches, its total cash transactional amount in 30 days can be divided by 50 to establish an adjusted risk factor and a risk score for a group comparison. Branches are used here as an example to measure the size of the operations. Other information, such as number of customers, number of transactions, number of employees, size of assets, etc. can also be used to measure the size of the operations.

[0108] In one aspect of the present disclosure, a set of risk factors, adjusted based on the size of operations (e.g., adjusted risk factors), are used to identify a group of entities which have this set of adjusted risk factors. The risk score of an adjusted risk factor is referred to as the adjusted risk score. If a particular entity has a total adjusted risk score that is much higher than a reference value derived from the total adjusted risk scores of all entities having the same set of adjusted risk factors, the particular entity may have conducted suspicious money laundering activities. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values. In general, in one aspect of the present disclosure, a detection algorithm that incorporates a risk factor in the detection algorithm can also be modified to incorporate an adjusted risk factor in the detection algorithm. A detection algorithm that incorporates a risk score in the detection algorithm can also be modified to incorporate an adjusted risk score in the detection algorithm.

[0109] To simplify the calculations, the standard group statistics such as mean, variance, standard deviation, etc., based on adjusted risk factors and adjusted risk scores can be derived to facilitate such comparison among a group of entities. As a result, even if there is no behavior change in

any account, a computer system can still detect suspicious money laundering activities based on the above approach.

[0110] Because MSBs may have different transactional activities from other types of businesses, it is more effective to monitor MSBs based on their unique transactional activities. Therefore, in one aspect of the present disclosure, a different set of detection algorithms may be used to monitor entities with a different set of risk factors. In one aspect of the present disclosure, a set of risk factors is used to identify a group of entities that have this set of risk factors, and a particular set of detection algorithms are used to detect suspicious money laundering activities in this group of entities. In other words, a set of detection algorithms is selected based on a set of risk factors associated with a group of entities to monitor the group of entities.

[0111] In another aspect of the present disclosure, a set of risk factors is adjusted based on the size of operations and used to identify a group of entities that have this set of adjusted risk factors, and a particular set of detection algorithms are used to detect suspicious money laundering activities in this group of entities. In other words, a set of detection algorithms is selected based on a set of adjusted risk factors associated with a group of entities to monitor the group of entities.

[0112] Sometimes, it is meaningful to monitor entities with higher risks more closely than entities with lower risks. Therefore, different sets of detection algorithms are used to monitor different entities with different levels of risk. In one aspect of the present disclosure, a set of detection algorithms is selected based a total risk score of an entity to monitor the entity. In another aspect of the present disclosure, a set of detection algorithms is selected based on a total adjusted risk score of an entity to monitor the entity where the total adjusted risk score is obtained from the risk scores of adjusted risk factors.

[0113] In one aspect of the present disclosure, once an MSB is detected as having possible money laundering activities, a computer system can identify the transaction (or a group of transactions) that caused the detected MSB to have a higher total risk score than a reference value derived from the total risk scores of all MSBs. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values.

[0114] Similarly, once an MSB is detected as having possible money laundering activities, a computer system identifies the transaction (or a group of transactions) that caused the detected MSB to have a higher total adjusted risk score than a reference value derived from the total adjusted risk scores of all MSBs. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values. As a result, a money laundering transaction (or a group of money laundering transactions) can be identified through this approach. This approach of identifying a particular transaction (or a group of transactions) with a higher risk score (or higher adjusted risk score) can be used for other types of customers, not just for MSBs.

[0115] Conventionally, a higher risk score implies a higher risk. There is no rule, however, that prohibits a person or business from defining a lower risk score for a higher risk. To avoid confusion, the description in this disclosure is based on the convention that a higher risk score means a

higher risk. Furthermore, a risk score can be a negative value. A negative risk score implies a reduced risk based on this convention.

[0116] As described above, an MSB is just one example. Other types of businesses, e.g., pawn shops, car dealers, etc., can be monitored in a similar way. As a result, risk factor, risk score, adjusted risk factor, adjusted risk score, total risk score, and total adjusted risk score can be used in a variety of methods to detect suspicious money laundering activities even if there is no behavior change in any account.

[0117] In fact, a government or non-government agency such as OCC, FDIC, FRB, NCUA, FinCEN, SEC, FINRA, etc., can monitor financial institutions such as banks, credit unions, insurance companies, stockbrokers, etc. based on similar approaches as described above for monitoring MSBs. Different risk factors, risk scores, adjusted risk factors and adjusted risk scores can be defined for this monitoring purpose.

[0118] In one aspect of the present disclosure, a computer system uses many different risk factors to determine whether a financial institution has complied with regulatory requirements for filing SARs to report money laundering and terrorist financing cases. For example, these risk factors may include a number of SARs filed for money laundering and terrorist financing cases, category of the financial institution, business type of the financial institution, geographical area of the financial institution, country of the head office of the financial institution, nature of the business of the financial institution, product types of the business, services types of the business, structure of the business, customer profile of the financial institution, historical records, type of the transactions conducted, funds inflows, funds outflows, transactional pattern, number of transactions, amount of transactions, transactional volume, transactional frequency, transactional derivatives, locations of the transaction, time of the transaction, countries of the transactions, senders of money transfer transactions, locations of the senders, countries of the senders, nature of the senders, recipients of money transfer transactions, locations of the recipients, countries of the recipients, natures of the recipients, relationships, social status of the customers, political exposure of the customers, political exposure of the senders, political exposure of the recipients, historical transactions, etc. In fact, thousands of risk factors may be considered to determine the compliance risk of a financial institution.

[0119] In one aspect of the present disclosure, the number of branches is used to adjust the risk factors and risk scores. In another aspect of the present disclosure, the asset size is used to adjust the risk factors and risk scores. Many other factors can also be used to adjust the risk factors and risk scores. In this current example, the "number of SARs filed" risk factor may have a negative value because the more SARs filed by the financial institution, the less chance that the financial institution has failed to file SARs.

[0120] In one aspect of the present disclosure, a set of risk factors is adjusted based on the size of operations and used to identify a group of banks that have this set of adjusted risk factors. If a particular bank has a total adjusted risk score that is much higher than a reference value of the total adjusted risk scores of all banks having the same set of adjusted risk factors, the particular bank may have not fulfilled its compliance obligation of detecting and reporting suspicious money laundering and/or terrorist financing activities. The reference value comprises an average, a

median, a mean, a mode, a weighted average, and/or other statistical values. To simplify the calculations, the standard group statistics such as mean, variance, standard deviation, etc. can be derived to facilitate such comparison among a group of entities.

[0121] Furthermore, different detection algorithms may be used to monitor different banks with different sets of risk factors. In one aspect of the present disclosure, a set of risk factors is used to identify a group of banks having this set of risk factors, and a particular set of detection algorithms is used to detect possible oversights in compliance matters in this group of banks. Therefore, in one aspect of the present disclosure, a set of detection algorithms is selected based on a set of risk factors associated with a group of banks for monitoring the group of banks.

[0122] In another aspect of the present disclosure, a set of risk factors is adjusted based on the size of operations and used to identify a group of banks that have this set of adjusted risk factors, and a particular set of detection algorithms are used to detect possible oversights in compliance matters in this group of banks. In other words, a set of detection algorithms is selected based on a set of adjusted risk factors associated with a group of banks to monitor the group of banks.

[0123] Although banks are used in the above example, the same set of methods can be used to monitor credit unions, stockbrokers, insurance companies, other financial institutions, and other types of businesses. Moreover, the monitoring scope is not limited to compliance with anti-money laundering and anti-terrorist financing matters. In fact, all types of matters of all types of businesses can be monitored by the methods described in the present disclosure by properly defining the risk factors, risk scores, adjusted risk factors, adjusted risk scores, and detection algorithms associated with such matters.

[0124] MSBs are also under pressure to comply with many laws and regulations. Unlike banks or credit unions, however, MSBs do not really know who their customers are. A typical MSB provides money services to any consumer who walks into its office. Even if the MSB collects identification information from all of its clients, an MSB may not be able to correctly identify money laundering activities. For example, it is possible that one consumer uses his Mexican Passport to conduct one money transfer transaction of \$7,000 by paying an MSB cash in the morning, and uses his California driver license to conduct another money transfer transaction of \$8,000 by paying the same MSB cash in the afternoon. Because two identification documents are used, this same consumer may be viewed as two different persons. The MSB may fail to file a Currency Transaction Report as required by law because more than \$10,000 cash has been provided by the same consumer. If the MSB has multiple branches, this situation becomes even more complicated because the same consumer can walk into different branches to conduct transactions based on different identification documents.

[0125] In one aspect of the present disclosure, a computer system compares the name, phone number, address, date of birth, etc. of all consumers who conducted transactions with an MSB to identify all transactions that may be conducted by the same consumer. After all transactions associated with a consumer have been identified, a computer system can

detect suspicious money laundering activities associated with the consumer based on the transactions associated with the consumer.

[0126] In one aspect of the present disclosure, a BSA Officer (e.g., a person tasked with investigating) investigates the detected case to determine whether it is a true money laundering case. The BSA Officer also reviews all historical cases associated with the consumer of the newly detected case. If the BSA Officer agrees that the detected case is a suspicious money laundering case, the computer system assists the BSA Officer to file a SAR with FinCEN. If the BSA Officer decides not to file a SAR, the BSA Officer enters a reason into the computer system to justify his decision for not reporting the detected case.

[0127] Sometimes, a bank receives a wire transfer from a client of corresponding bank A and resends the wire transfer to another client of corresponding bank B because corresponding bank A and corresponding bank B do not have a direct banking relationship. This situation often arises during international wire transfers because banks in two different countries may not have a direct banking relationship. This type of wire transfer is often referred to as intermediary wire transfer.

[0128] A bank that provides intermediary wire transfer services is exposed to a very high money laundering risk because senders and recipients of intermediary wire transfers are not customers of the bank. In addition, the bank may not know the real background of the senders and the recipients of the wire transfers. It is possible that a sender is a terrorist financier and a recipient is the terrorist. A bank that handles intermediary wire services may unknowingly become a channel for money laundering and terrorist financing.

[0129] In one configuration of the present disclosure, a computer system compares names, addresses, countries, phone numbers, email addresses, etc. of all senders and recipients of intermediary wire transfers and identifies the transactions associated with each sender and each recipient. In one aspect of the present disclosure, if the computer system detects an unusually large number of wire transfers from the same sender, the sender and the recipients may be involved in money laundering or terrorist financing activities. If the computer system detects an unusually large aggregate amount of wire transfers from the same sender, the sender and the recipients may be involved in money laundering activities.

[0130] Similarly, if the computer system detects an unusually large number of wire transfers to the same recipient, the senders and the recipient may be involved in money laundering or terrorist financing activities. If the computer system detects an unusually large aggregate amount of wire transfers to the same recipient, the senders and the recipient may be involved in money laundering activities.

[0131] If the computer system detects that an unusual number of wire transfers are sent from the same sender to the same recipient, the sender and the recipient may be involved in money laundering or terrorist financing activities. If the computer system detects that an unusual aggregate amount of wire transfers is sent from the same sender to the same recipient, the sender and the recipient may be involved in money laundering or terrorist financing activities.

[0132] In one aspect of the present disclosure, a BSA Officer investigates such a detected case to determine whether it is a true money laundering case. The BSA Officer

also reviews all historical cases associated with the suspects of the newly detected case. If the BSA Officer agrees that there is suspicious money laundering activity, the computer system assists the BSA Officer to file a SAR to the FinCEN. If the BSA Officer decides not to file a SAR, the BSA Officer enters a reason into the computer system to justify his decision for not reporting such detected activities.

[0133] As a large percentage of the population is quickly aging, the Elder Abuse Reporting Act (EARA) was recently established in some states to protect elders who are unable to protect themselves. Very often, an elderly person may give money to a perpetrator because the elderly person is cheated by the perpetrator. Therefore, financial institutions are training frontline personnel to observe and report what they see as possible elder abuse cases. This human-based approach is not effective because transactions can be performed remotely and perpetrators can skillfully cover up their activities. Moreover, human workers are susceptible to errors and mistakes. It is ineffective to depend on human workers to detect and report elder abuse cases.

[0134] For many businesses, the date of birth information of their customers is stored inside a database. In one aspect of the present disclosure, a computer system collects the date of birth information and identifies elderly persons who are older than a pre-defined age. A computer system monitors the transactions of all elderly persons and detects any change in the activities of these elderly persons.

[0135] For example, if an unusually large amount of funds is sent out from an account of an elderly person, the financial institution may want to investigate the purpose of the funds transfer. In one aspect of the present disclosure, if a check with an unusually large amount is deposited into an account of an elderly person, the financial institution may want to investigate whether a fake check was given to the elderly person in exchange for his/her real money or assets. If there is an unusual transactional pattern, (e.g., unusual frequency or volume) in the account of an elderly person, the financial institution may want to investigate the transaction(s). If the account balance of an elderly person is decreasing rapidly, the financial institution may want to investigate the transactions associated with this account.

[0136] In one aspect of the present disclosure, risk factors, risk scores, adjusted risk factors, adjusted risk scores, total risk scores, total adjusted risk score, statistical approach, and the methods to select detection algorithms described earlier can be applied to detect possible elder abuse cases. Because elder abuse is different from money laundering, a different set of risk factors and risk scores may be used for elder abuse detection. For example, these risk factors may include age of the person, sex of the person, income level of the person, appearance of the person, judgment about the person, personal conditions of the person, family conditions of the person, family members of the person, family members' conditions of the person, friends of the person, friends' conditions of the persons, historical records of the person, industry category of the person, geographical area of the person, country of the address of the person, profession of the person, nationality, type of the transaction conducted, balance of the accounts, funds inflow, funds outflow, transactional pattern, number of transactions, amount of transactions, transactional volume, transactional frequency, transactional derivative, location of the transaction, time of the transaction, country of the transaction, sender of a money transfer transaction, location of the sender, country of the

sender, nature of the sender, recipient of a money transfer transaction, location of the recipient, country of the recipient, nature of the recipient, relationships, social status, political exposure, historical transactions, etc. In fact, many different risk factors may be considered to determine the elder abuse risk of a person.

[0137] For example, in one aspect of the present disclosure, a risk factor is used to identify a group of elderly persons having the same risk factor. If a particular elderly person has a total risk score higher than a reference value derived from the total risk scores of all elderly persons having the same risk factor, the particular elderly person may become a victim of a potential elder abuse case. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values. In another aspect of the present disclosure, a set of risk factors is used to identify a group of elderly persons who have this set of risk factors. If a particular elderly person has a total risk score higher than a reference value derived from the total risk scores of all elderly persons having the same set of risk factors, the particular elderly person may become a victim of a potential elder abuse case. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values.

[0138] To simplify the calculations, the standard group statistical values such as mean, variance, standard deviation, etc. can be derived to facilitate such comparison among a group of entities. As a result, even if there is no behavior change in an account, a computer system can still detect a possible elder abuse case based on the above approach.

[0139] Very often, a business may have a Compliance Officer that is responsible for all regulatory compliance matters. In one aspect of the present disclosure, an investigator (e.g., a Compliance Officer) investigates detected cases to determine whether true elder abuse cases have occurred. The Compliance Officer also reviews all historical cases associated with the elderly person of the newly detected case. If the Compliance Officer agrees that the case is a possible elder abuse case, the computer system assists the Compliance Officer in reporting the detected case. If the Compliance Officer decides not to report the detected case, the Compliance Officer enters a reason into the computer system to justify his decision for not reporting the detected case.

[0140] According to the Sarbanes-Oxley Act (SOX), certain companies (e.g., publicly listed companies) have to conduct internal control monitoring to prevent fraud committed by employees. Traditionally, such internal control monitoring is performed by human workers (e.g., auditors) that spend several months each year auditing the financial records of a business. Such a human-based approach is not effective because human workers are prone to errors and mistakes. Furthermore, because it takes a substantial amount of time to audit financial records, it may be too late to prevent a crime.

[0141] In one aspect of the present disclosure, a computer system monitors accounting general ledger items and detects any unusual patterns (e.g., unusual frequency, volume, acceleration, etc.) related to the general ledger items to identify suspicious internal fraud activities. For example, if a traveling expense general ledger item has suddenly grown by 500% this month when it is compared with the past twelve-month history, some employees may have abused their rights and caused the unusual expenses.

[0142] In one aspect of the present disclosure, a computer system compares the current value of an accounting general ledger item with a reference value derived from the historical values of the same accounting general ledger item for the past x months, where the value x is predefined. If the current value is larger than the reference value by a noticeable margin, some employee may have committed fraud. The reference comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values. Further investigation can be conducted to determine why the general ledger item value has deviated from its historical values.

[0143] In another aspect of the present disclosure, a computer system compares an employee's current activities with his historical activities to detect any change. For example, if a loan officer has issued an unusually larger number of loans per month in comparison to historical monthly numbers, this loan officer's activities may be suspicious. If a loan officer has issued a loan with a loan amount unusually larger than the historical amounts, this loan officer's activities may be suspicious. If a loan officer has issued an unusual larger aggregate amount of loans per month compared to the historical monthly aggregate amounts, this loan officer's activities may be suspicious.

[0144] Very often, an activity can be measured by a value referred to as an activity value. For example, a loan officer's activities can be measured by the number of loans, largest amount of loan, aggregate amount of loans, average amount per loan, number of loans to the same customer, number of changes on loan records, number of changes on loan records with the same customer, frequency of changes on loan records, frequency of changes on loan records with the same customer, types of loans, etc. A bank teller's activities can be measured by the total number of transactions, aggregate amount of transactions, largest amount of transaction, average amount per transaction, types of transactions, the number of customers transacting business with the teller, the average number of transactions per customer, the number of transactions with the same customers, number of changes on customer records, number of changes on customer records with the same customer, frequency of changes on customer records, frequency of changes on customer records with the same customer, etc. In one aspect of the present disclosure, a computer system compares the current value of an activity with a reference value derived from the historical values of the same activity. When the current value is larger than the reference value by a noticeable margin, the person who conducted the activity may have committed fraud. Further investigation can be conducted to determine whether this person has truly committed fraud. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values.

[0145] In one aspect of the present disclosure, a computer system compares an employee's activity with the activities of other employees who have the same role in the business. For example, if one teller (or loan officer, etc.) behaves very differently from other tellers (or loan officers, etc.) in the same branch, this teller (or loan officer, etc.) may have conducted some suspicious activities.

[0146] In one aspect of the present disclosure, a computer system compares an activity value of a particular employee with a reference value derived from all the activity values for the same activity of all employees who have the same responsibilities as the particular employee. When the activity value of the particular employee noticeably deviates from

the reference value, the particular employee may have committed fraud. Further investigation can be conducted to determine whether this employee has truly committed fraud. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values.

[0147] When one employee is compared with a group of employees, the statistical approach used in the flight attendant example described earlier can be applied. For example, a comprehensive set of risk factors associated with employees can be identified and a risk score is assigned to each risk factor. As a result, each employee has a total risk score that is obtained from a mathematical transformation (e.g., summation) of all risk scores associated with the employee.

[0148] The set of risk factors for detecting fraud associated with employees may be different from the set of risk factors to detect other type of suspicious activities, such as money laundering. For example, the risk factors for detecting employee fraud may include job type of the employee, education level of the employee, income level of the employee, length of employment at the current job, performance review records, employment history, duration of each employment in the employment history, reasons for termination of each employment in the employment history, age of the employee, sex of the employee, personal conditions of the employee, family conditions of the employee, family members of the employee, family members' conditions of the employee, friends' conditions of the employees, historical records of the employee, type of work performed, number of transactions performed, amount of transactions performed, largest amount of transaction, number of transactions with a particular counter party, amount of transactions with a particular counter party, number of changes of crucial records, number of changes of crucial records associated with a particular counter party, geographical area of the employee's home, geographical area of the employee's office, country of the address of the employee, nationality, type of the transaction conducted, balance of the accounts, funds inflow, funds outflow, transactional pattern, number of transactions, amount of transactions, transactional volume, transactional frequency, transactional derivative, location of the transaction, time of the transaction, country of the transaction, sender of a money transfer transaction, location of the sender, country of the sender, nature of the sender, recipient of a money transfer transaction, location of the recipient, country of the recipient, nature of the recipient, relationships, social status, political exposure, historical transactions, etc. In fact, numerous risk factors may be considered to determine the employee fraud risk. In one aspect of the present disclosure, different sets of risk factors may be used to detect different types of suspicious activities.

[0149] In one aspect of the present disclosure, when the total risk score of a particular employee is higher than the mean of total risk scores of all employees having the same risk factor as the particular employee has by a noticeable margin, this particular employee may have conducted suspicious activities. The noticeable margin can be set in terms of a number of the standard deviations or other reference values.

[0150] Instead of one risk factor, multiple risk factors can be used to improve the accuracy of the detection result. In one aspect of the present disclosure, if the total risk score of a particular employee is higher than the mean of total risk scores of all employees having the same set of risk factors

as the particular employee has by a noticeable margin, this particular employee may have conducted some suspicious activities. In one example, the noticeable margin is set in terms of a number of the standard deviations or other reference values.

[0151] In fact, by identifying the risk factors associated with a group of entities and properly assigning a risk score to each risk factor, a statistical approach based on the total risk score of each entity for identifying a suspicious activity of a particular entity can be applied to many other situations in addition to money laundering, terrorist financing, and employee fraud.

[0152] In one aspect of the present disclosure, many risk factors are associated with a group of entities. Each of the risk factors can be assigned a risk score. Each entity can be given a total risk score based on a mathematical transformation, such as a summation. For example, other possible mathematical transformations include, but are not limited to multiplication, division, and subtraction, summation of squares, square of summation, a mix of the above, and other like ways of combining risk scores.

[0153] In one aspect of the present disclosure, when the total risk score of a particular entity is higher than the mean of total risk scores of all entities having the same risk factor as the particular entity has by a predefined margin, this particular entity may have conducted some suspicious activities. The predefined margin can be set in terms of a number of the standard deviations or other reference values.

[0154] In another aspect of the present disclosure, if the total risk score of a particular entity is higher than the mean of total risk scores of all entities having the same set of risk factors as the particular entity has by a predefined margin, this particular entity may have conducted some suspicious activities

[0155] In one aspect of the present disclosure, a computer system identifies one transaction (or a group of transactions) that has caused the particular entity to have a total risk score that is higher than the mean of total risk scores of all entities. Such a transaction (or a group of transactions) may be a suspicious activity.

[0156] The noted statistical approach is just one way of managing risks. Many other group comparison methods can also be used. Furthermore, suspicious activities may not be limited to illegal or prohibited activities. An activity becomes suspicious because it is different from normal activities. It may be harmless or even possibly an activity with good intentions. Therefore, investigation is often required to make the final determination of whether to report the detected case.

[0157] In one aspect of the present disclosure, a responsible person investigates a newly detected case to determine whether it is illegal. The responsible person also reviews all historical cases associated with the suspect(s) of the newly detected case. When the responsible person agrees that the detected case is illegal, the computer system assists the responsible person to report the detected case. When the responsible person decides not to report the detected case, the responsible person enters a reason into the computer system to justify his/her decision for not reporting the detected case.

[0158] After the 9/11 tragedy, the U.S. Congress passed the Unlawful Internet Gambling Enforcement Act (UIGEA) because online gambling can be a vehicle for conducting money laundering and terrorist financing activities. Regula-

tion GG was established in response to the Unlawful Internet Gambling Enforcement Act. According to Regulation GG, a financial institution needs to ask questions during the account opening process as to whether a new customer will conduct any online gambling activity. Because perpetrators know online gambling is illegal, they will lie during the account opening process. As a result, the "question asking" approach defined in Regulation GG is simply a formality. However, Regulation GG has specifically stated that Regulation GG does not modify the obligation of a financial institution to file a SAR under the Bank Secrecy Act.

[0159] In other words, if a perpetrator lied during the account opening process and actually conducts an illegal online gambling business, the financial institution has the obligation to report the case to FinCEN through a SAR. In one aspect of the present disclosure, a computer system compares the senders and recipients of all funds transfer transactions during a period of time. If a customer has sent a large amount of money to a recipient and has also received a large amount of money from the same recipient during a period of time, such transactions may possibly be deposits of wager funds and payments for money earned from gambling activities between an online gambler and an online gambling organization. The computer system detects such cases as possible illegal online gambling cases. Once a case is detected, further investigation is required.

[0160] In one aspect of the present disclosure, a computer system detects a customer as a possible online gambling organization when the computer system detects a large number of transactions with large dollar amounts associated with the customer because an online gambling organization typically deals with a large amount of money and a large number of clients. The computer system detects such a case as a possible illegal online gambling case. Once a case is detected, further investigation is required.

[0161] In one aspect of the present disclosure, a computer system compares a list of known names of online gambling organizations with the senders and the recipients of funds transfer transactions associated with a customer. If there is a match, the customer may be involved in online gambling activities. The computer system detects this case as a possible illegal online gambling case. Once a case is detected, further investigation is required.

[0162] In addition to the noted transactional pattern monitoring, group comparison methods described earlier can also be applied to detect possible illegal online gambling activities. In one aspect of the present disclosure, all risk factors related to online gambling are identified. For example, these risk factors may include due diligence results of the customer, length of the account history, industry category of the customer, business type of the customer, number of name matches with gambling organizations in transactions, geographical area of the customer, country of the head office of the customer, nature of the business of the customer, product types of the business, services types of the business, structure of the business, profession of the customer, nationality, historical records, type of the transaction conducted, balance of the accounts, funds inflow, funds outflow, transactional pattern, number of transactions, amount of transactions, transactional volume, transactional frequency, transactional derivative, number of charge backs, location of the transaction, time of the transaction, country of the transaction, sender of a money transfer transaction, location of the sender, country of the sender, nature of the sender, recipient

of a money transfer transaction, location of the recipient, country of the recipient, nature of the recipient, relationships, social status, political exposure, historical transactions, etc. In fact, many different risk factors may be considered to determine online gambling risk. As explained earlier in this disclosure, adjusted risk factors can also be used so that adjusted risk scores can be applied based on the size of operations.

[0163] In one aspect of the present disclosure, a risk factor is used to identify a group of customers that have the same risk factor. When a particular customer has a total risk score higher than a reference value derived from the total risk scores of all customers with the same risk factor, the particular customer may be involved in illegal online gambling. In another aspect of the present disclosure, a set of risk factors is used to identify a group of customers who have this set of risk factors. If a particular customer has a total risk score higher than a reference value derived from the total risk scores of all customers with the same set of risk factors, the particular customer may be involved in illegal online gambling. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values. To simplify the calculations, the standard group statistics such as mean, variance, standard deviation, etc. can be derived to facilitate the comparison among a group of customers.

[0164] In one aspect of the present disclosure, a responsible person (or the BSA Officer) investigates the detected case to determine whether it is a true online gambling case. The BSA Officer also reviews all historical cases associated with the suspect of the newly detected case. When the BSA Officer agrees that the detected case is a possible illegal online gambling case, the computer system assists the BSA Officer in filing a SAR to the FinCEN. When the BSA Officer decides not to file a SAR, the BSA Officer enters a reason into the computer system to justify his decision for not reporting the detected case.

[0165] The U.S. Congress has passed the Fair and Accurate Credit Transactions Act (FACT Act) to protect consumers. In particular, businesses are expected to identify and report identity theft cases. Financial institutions are also expected to file a SAR when identity theft cases are detected.

[0166] In one aspect of the present disclosure, a computer system monitors consumer reports and other available information to detect a fraud or active duty alert that is included in a consumer report, a notice of a credit freeze, and/or a notice of an address discrepancy. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0167] In one aspect of the present disclosure, a computer system monitors consumer reports and available information to detect a consumer report indicating a pattern of activity that is inconsistent with the history and usual pattern of activity of an applicant or customer. For example, a recent and significant increase in the volume of inquiries, an unusual number of recently established credit relationships, a material change in the use of credit, especially with respect to recently established credit relationships, or an account that was closed for cause or identified for abuse of account privileges by a financial institution or creditor may represent an unusual pattern. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0168] In one aspect of the present disclosure, a computer system detects whether documents provided for identification appear to have been altered or forged. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0169] In one aspect of the present disclosure, a computer system detects whether the photograph or physical description on the identification is not consistent with the appearance of the applicant or customer presenting the identification. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0170] In one aspect of the present disclosure, a computer system detects whether other information on the identification is inconsistent with information provided by the person opening a new account or presenting the identification. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0171] In one aspect of the present disclosure, a computer system detects whether other information on the identification is inconsistent with readily accessible information that is on file with the financial institution or creditor, such as a signature card or a recent check. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0172] In one aspect of the present disclosure, a computer system detects whether an application appears to have been altered or forged, or gives the appearance of having been destroyed and reassembled. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0173] In one aspect of the present disclosure, a computer system determines whether personal identifying information provided is inconsistent when compared against external information sources used by the financial institution or creditor. For example, the address may not match any address in the consumer report, or the Social Security Number (SSN) has not been issued, or is listed on the Social Security Administration's Death Master File. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0174] In one aspect of the present disclosure, a computer system determines whether some personal identifying information provided by the customer is inconsistent with other personal identifying information provided by the customer. For example, there may be a lack of correlation between the SSN range and date of birth. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0175] In one aspect of the present disclosure, a computer system determines whether personal identifying information provided is associated with known fraudulent activity as indicated by internal or third-party sources used by the financial institution or creditor. For example: the address on an application may be the same as the address provided on a fraudulent application; or the phone number on an application may be the same as the number provided on a fraudulent application. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0176] In one aspect of the present disclosure, a computer system determines whether personal identifying information provided is of a type commonly associated with fraudulent

activity as indicated by internal or third-party sources used by the financial institution or creditor. For example, the address on an application may be fictitious, a mail drop, or a prison; or the phone number may be invalid, or associated with a pager or answering service. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0177] In one aspect of the present disclosure, a computer system determines whether the Social Security Number provided is the same as or similar to the account number or telephone number submitted by an unusually large number of other persons opening accounts or other customers. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0178] In one aspect of the present disclosure, a computer system determines whether the address or telephone number provided is the same as or similar to the account number or telephone number submitted by an unusually large number of other persons opening accounts or other customers. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0179] In one aspect of the present disclosure, a computer system determines whether the person opening the account fails to provide all required personal identifying information on an application or in response to notification that the application is incomplete. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0180] In one aspect of the present disclosure, a computer system determines whether personal identifying information provided is inconsistent with personal identifying information that is on file with the financial institution or creditor. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0181] In one aspect of the present disclosure, a computer system determines whether the person opening the account cannot provide authenticating information such as answers to challenge questions, which are beyond that which generally would be available from a wallet or consumer report. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0182] In one aspect of the present disclosure, a computer system determines whether there is unusual use of, or suspicious activity related to, the account. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0183] In one aspect of the present disclosure, a computer system determines whether, shortly following the notice of a change of address for an account, the institution or creditor receives a request for a new, additional, or replacement card or a cell phone, or for the addition of authorized users on the account. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0184] In one aspect of the present disclosure, a computer system determines whether a new revolving credit account is used in a manner commonly associated with known patterns of fraud. For example: the majority of available credit is used for cash advances or merchandise that is easily convertible to cash (e.g., electronics equipment or jewelry); or the customer fails to make the first payment or makes an initial payment but no subsequent payments. If a suspicious

activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0185] In one aspect of the present disclosure, a computer system determines whether an account is used in a manner that is inconsistent with established patterns of activity on the account. There is, for example: nonpayment when there is no history of late or missed payments, a material increase in the use of available credit, a material change in purchasing or spending patterns, a material change in electronic fund transfer patterns in connection with a deposit account; or a material change in telephone call patterns in connection with a cellular phone account. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0186] In one aspect of the present disclosure, a computer system determines whether an account that has been inactive for a reasonably lengthy period of time is used (taking into consideration the type of account, the expected pattern of usage and other relevant factors). If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0187] In one aspect of the present disclosure, a computer system determines whether mail sent to the customer is returned repeatedly as undeliverable although transactions continue to be conducted in connection with the customer's account. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0188] In one aspect of the present disclosure, a computer system closely reviews all transactions when the financial institution or creditor is notified that the customer is not receiving paper account statements. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0189] In one aspect of the present disclosure, a computer system closely reviews all transactions when the financial institution or creditor is notified of unauthorized charges or transactions in connection with a customer's account. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0190] In one aspect of the present disclosure, a computer system closely reviews all transactions when the financial institution or creditor is notified by a customer, a victim of identity theft, a law enforcement authority, or any other person that it has opened a fraudulent account for a person engaged in identity theft. If a suspicious activity case is detected, the computer system makes the detected case available for a responsible person to review.

[0191] In addition to monitoring transactional patterns as described above, group comparison methods described earlier can also be applied to detect possible identity theft cases. Identity theft cases can be classified into two primary categories. The first category includes cases in which a victim's account, financial instrument, or identification document is stolen by a fraudster to conduct activities. Under such circumstances, as described earlier, a computer system can detect activities that deviate from the expected activities of the victim, which can be established from the historical activities of the victim.

[0192] The second category includes cases in which a victim's identity is stolen to open a new account and/or start some new activities. Under such circumstances, the victim is out of the picture from day one. Because there is no true

historical activity of the victim, the expected activities of the victim cannot be correctly established for fraud prevention purposes. Although somebody can ask the perpetrator some questions and collect answers during the account opening process with the intention of establishing expected activities of the perpetrator, this question-and-answer approach may not work because the perpetrator knows how to answer questions for establishing his expected activities without triggering any alert.

[0193] To detect identity theft when there is no true historical activity available, in one aspect of the present disclosure, all risk factors of a new account or new customer are identified. For example, these risk factors may include due diligence results of the customer, prior records of the customer with other businesses, credit report records of the customer, industry category of the customer, business type of the customer, geographical area of the customer, country of the address of the customer, nature of the business of the customer, product types of the business, services types of the business, structure of the business, profession of the customer, nationality, historical records, type of the transaction conducted, balance of the accounts, funds inflow, funds outflow, transactional pattern, number of transactions, amount of transactions, transactional volume, transactional frequency, transactional derivative, number of charge backs, location of the transaction, time of the transaction, country of the transaction, sender of a money transfer transaction, location of the sender, country of the sender, nature of the sender, recipient of a money transfer transaction, location of the recipient, country of the recipient, nature of the recipient, relationships, social status, political exposure, historical transactions, etc. In fact, numerous risk factors may be considered to determine the identity theft risk.

[0194] In one aspect of the present disclosure, a risk factor is used to identify a group of persons who have the same risk factor. If a particular person has a total risk score which is much higher than a reference value derived from the total risk scores of all persons who have the same risk factor, the particular person may be involved in an identity theft case. A set of risk factors can be used to identify a group of persons who have this set of risk factors. If a particular person has a total risk score higher than a reference value derived from the total risk scores of all persons with the same set of risk factors, the particular person may be involved in an identity theft case. The reference value comprises an average, a median, a mean, a mode, a weighted average, and/or other statistical values. To simplify the calculations, group statistics such as mean, variance, standard deviation, etc. can be derived to facilitate such comparison among a group of persons.

[0195] In one aspect of the present disclosure, a responsible person (or the Compliance Officer) investigates the detected case to determine whether it is a true identity theft case. The Compliance Officer also reviews all historical cases associated with the newly detected case. If the Compliance Officer agrees that the case is a possible identity theft case, the computer system assists the Compliance Officer to file a SAR with the FinCEN. If the Compliance Officer decides not to file a SAR, the Compliance Officer enters a reason into the computer system to justify his decision for not reporting the detected activities.

[0196] The Office of Foreign Assets Control (OFAC) has a very simple rule stating that it is illegal to have any business transaction with any entity on the list published by

the Office of Foreign Assets Control. This list is commonly referred to as the "OFAC List." This rule applies to all U.S. persons and entities, including financial institutions. For example, Wal-Mart was fined by the OFAC for a violation of this rule. U.S. financial institutions that are under the tightest regulatory monitoring naturally must strictly comply with this rule.

[0197] In the beginning, it was a very simple rule. The implications of this rule, however, have become much more complicated over the past 20 years. A common issue occurs when people misspell (including mistype, mispronounce, etc.) their names. Even if an entity's name was misspelled but is on the OFAC List, a financial institution still has the obligation of identifying this entity as an entity on the OFAC List (generally referred to as an OFAC match).

[0198] The natural question is how much deviation from the original names on the OFAC List is classified as "misspelling." The OFAC and government regulators have never given any precise guidance for answering this question. A very common exercise an examiner or auditor may perform is the use a notorious name like "Osama bin Laden" as a sample to test a business. In general, a business is supposed to identify all business transaction associated with "Osama bin Laden," "Osama Laden," "Osama Laten," "Laten Osama," "Latin Obama," etc. as possible OFAC matches. Now, if the scope of deviation from the OFAC names is further broadened, whether a financial institution is supposed to identify a single word "Obama" which is a former U.S. president's name as a possible OFAC match is questionable. One can easily tell that such a simple OFAC rule has created a great deal of confusions in recent years.

[0199] In one aspect of the present disclosure, an "OFAC Match Scale" is used to measure the degree of deviation. A value referred to as "Relative Correlation" ("RC value") may be generated by the OFAC Match Scale to measure the similarity between two names. For example, if a name has an RC value of 100%, it exactly matches an OFAC name on the OFAC list. If a name has an RC value of 97%, it may be different from an OFAC name on the OFAC list by one letter or two letters. If a name has an RC value of 0%, it is totally different from all OFAC names on the OFAC list.

[0200] In one aspect of the present disclosure, the length of the name also affects the RC value. For example, if a name is different by one letter from an OFAC name, which has 25 letters, the RC value may be 96%, while another name may have an RC value of 90% although it is also different by just one letter from another OFAC name, which has 10 letters.

[0201] Some long words, such as international, incorporation, limited, company, organization, etc., are commonly used for business names, and such words also exist in the OFAC names list. As a result, these long words generate higher RC values for those businesses that use these long words in their names. To avoid unnecessary false positives, in one aspect of the present disclosure, a commonly used long word may be replaced with a short word to reduce its impact on the RC value. For example, the word "international" can be replaced with "intl."

[0202] In addition, some countries do not use the descriptions "first name" and "last name." As a result, a person may use a different sequence of names when this person is asked to provide a first name and a last name. "Osama Laden" may

become “Laden Osama.” In one aspect of the present disclosure, an OFAC match scale identifies a possible “off-sequence” OFAC match.

[0203] Moreover, some words are commonly used in certain cultures without contributing distinct differentiations. For example, “bin” means “son of” and “binti” means “daughter of” in the Muslim culture. Official names in the Muslim culture have either “bin” or “binti” in the name. For example, if a Muslim father has the name “John,” his daughter “Mary” will have an official name of “Mary binti John” and his son “David” will have an official name “David bin John.” Under such circumstances, the words “bin” and “binti”, which are commonly used in Muslim names, will create “false similarity” between two Muslim names. To provide a more scientifically correct result, in one aspect of the present disclosure, the OFAC match scale may exclude these kinds of “trivial words” before the calculation of the RC value. Sometimes, a name may be translated into English based on sound. Therefore, in one aspect of the present disclosure, the OFAC match scale should measure the match of sound to determine the RC value.

[0204] In one aspect of the present disclosure, a financial institution decides what threshold to use when it conducts an OFAC check. For example, if a financial institution uses a threshold of 75%, a possible OFAC match is detected when a name has an RC value of 75% or higher. Because each financial institution may have a different risk exposure from others, it is very likely that X is the best threshold for financial institution A while Y is the best threshold for financial institution B. As a general guideline, the X or Y value is selected in accordance with a risk-based principle.

[0205] In general, the higher threshold a financial institution uses, the fewer possible OFAC matches the financial institution detects. This saves time during the review process because more false positives are avoided. If the threshold is too high, however, it is possible that the financial institution may miss a reasonable deviation from an OFAC name, such as “Osama bin Laden.” If the threshold is too low, it is possible that the financial institution may falsely detect many of its clients as possible OFAC matches. The best practice is to find a tradeoff between “too many possible OFAC matches to review” and “missing real OFAC name deviations caused by misspellings.”

[0206] In one aspect of the present disclosure, a user can randomly select a number of OFAC names from the OFAC List and find out how the OFAC match scale responds to deviations from these selected OFAC names. A user can then determine when he calls a “possible OFAC match” based on this test. It is advisable to keep this test result for future auditors and examiners to review.

[0207] There is a possibility that a particular name is very close to an OFAC name. For example, American Express, a very reputable credit card company, is often falsely detected as an OFAC match because of the word “express.” Therefore, to avoid this type of frequent false positive, in one aspect of the present disclosure, an exempt list is generated by users to include those well-known reputable businesses on the exempt list. The businesses on the exempt list are classified as false positives, either automatically by the computer or manually by users, when they are detected as possible OFAC matches.

[0208] Very often, a business may have an OFAC officer that handles all OFAC related matters. In one aspect of the present disclosure, if an OFAC Officer (e.g., a responsible

person) of a financial institution detects a possible OFAC match with an RC value over a pre-defined threshold, the OFAC Officer investigates whether this is a true OFAC match. If the OFAC Officer believes that it is a true match, the OFAC Officer should handle the case according to the guidelines issued by the Office of Foreign Assets Control. According to the OFAC regulations, in some cases, the OFAC Officer may need to block the transaction so that the person on the OFAC list does not benefit from the transaction. If the OFAC Officer decides that the OFAC match is a false positive after his investigation, the OFAC Officer should enter a reason into the computer system to justify his reason for not reporting such OFAC match case to the Office of Foreign Assets Control and/or not block the transaction.

[0209] Section 314(a) of the USA PATRIOT Act requires financial institutions to detect matches of names on a 314(a) list, which is periodically published by the FinCEN. The computer system can handle the 314(a) compliance matters using a similar approach to the handling of the OFAC compliance matters, as described above.

[0210] Sometimes, a 314(a) list also includes additional personal identification information, such as identification document number, date of birth, address, etc. In one aspect of the present disclosure, in addition to the methods described above for detecting possible OFAC matches, personal identification information, such as identification document number, address, and/or date of birth, etc. is used by the computer system to determine whether a detected 314(a) match is a true match. This approach can reduce false positives in the 314(a) matching process.

[0211] In one aspect of the present disclosure, if a Compliance Officer (e.g., a responsible person) of a financial institution detects a possible 314(a) match with an RC value over a pre-defined threshold, the Compliance Officer investigates whether this is a true 314(a) match. In one aspect of the present disclosure, if the Compliance Officer believes it is a true match, the Compliance Officer reports the 314(a) match case to the FinCEN. If the Compliance Officer decides that the 314(a) match is a false positive after his investigation, the Compliance Officer enters a reason into the computer system to justify his reason for not reporting the 314(a) match to the FinCEN.

[0212] In one aspect of the present disclosure, the computer system receives customer information and transactional data from the core data processing system of the financial institution or other data processing systems that may be internal to or external from the financial institution. The customer information may include background information.

[0213] In one aspect of the present disclosure, the computer system receives information regarding suspicious activities observed by frontline personnel. For example, the computer system may receive information input from the frontline personnel. The computer system may also receive information provided by other internal or external sources.

[0214] Although “financial institutions” are used as examples for easy explanation, the present disclosure also applies to other types of businesses. In general, any business that needs to comply with laws and regulations may employ an intelligent alert system as described in the present disclosure.

[0215] In one aspect of the present disclosure, the risk score of a risk factor or a degree of a risk factor can be assigned by a computer software module, a person who

designs or tunes the system, or a user who uses the system. In most cases, the absolute value of the risk score is not crucial and the relative relationships among all risk scores may be more important.

[0216] Moreover, a subject's total risk score should only fluctuate within a reasonable range. In one aspect of the present disclosure, if a subject's total risk score has suddenly increased and is beyond a threshold, the subject may have conducted a suspicious or unusual activity. That is, if a difference between a subject's first total risk score and second total risk score is greater than an increase threshold, where the first total risk score is less than the second total risk score, the subject may have conducted a suspicious or unusual activity. In another aspect of the present disclosure, if a subject's total risk score has suddenly decreased tremendously, the subject may also have conducted a suspicious or unusual activity. That is, if a difference between a subject's second total risk score and first total risk score is greater than a decrease threshold, where the first total risk score is greater than the second total risk score, the subject may have conducted a suspicious or unusual activity. Therefore, an alert will be sent to an investigator, a BSA Officer, a Compliance Officer, or another type of responsible people who will investigate the subject when its total risk score has suddenly increased or decreased tremendously.

[0217] The observed data of a subject may fluctuate from time to time. Thus, the intelligent alert system may permit a certain range of fluctuation of the subject's total risk score to avoid false alerts. In one aspect of the present disclosure, the intelligent alert system increases a subject's permissible total risk score fluctuation range when the subject's total risk score is lower than a threshold. In another aspect of the present disclosure, the intelligent alert system decreases a subject's permissible total risk score fluctuation range when the subject's total risk score is higher than a threshold. The permitted fluctuation range may be determined (e.g., set) by a software module, a person who designs the system, a person who tunes the system, or a person who uses the system.

[0218] For example, if a subject's total risk score is higher than the mean of the total risk scores of all subjects plus a certain number of standard deviations of all risk scores, such as four standard deviations, the intelligent alert system may modify the subject's permissible total risk score fluctuate range to be within half a standard deviation without triggering an alert. In another example, if a subject's total risk score is within the mean of the total risk scores of all subjects plus a certain number of standard deviations, such as three standard deviations, the intelligent alert system may permit the subject's total risk score to fluctuate within a range of one standard deviation without triggering an alert.

[0219] In yet another example, if a subject's total risk score is within the mean of the total risk scores of all subjects plus a certain number of standard deviations, such as two standard deviations, the intelligent alert system may permit the subject's total risk score to fluctuate within a range of one and a half standard deviations without triggering an alert. In still yet another example, if a subject's total risk score is within the mean of the total risk scores of all subjects plus a certain number of standard deviations, such as one standard deviation, the intelligent alert system may permit the subject's total risk score to fluctuate within a range of two standard deviations without triggering an alert.

[0220] In the field of machine learning, a negative is a set of data that has not triggered an alert. A true negative is a set of data that has not triggered the alert and does not comprise a true case for triggering the alert. A false negative is a set of data that has not triggered the alert but comprises a true case for triggering the alert that the system has missed. As an example, a false negative money laundering case may cause a financial institution to be penalized by the U.S. government if the false negative case is discovered by the U.S. government. Therefore, it is desirable to prevent false negatives in an alert system specified to prevent money laundering (e.g., anti-money laundering alert system).

[0221] For an anti-money laundering alert system in a U.S. financial institution, a true money laundering case is reported to FinCEN, a U.S. government organization. FinCEN has a set of communication protocols. Financial institutions in the U.S. can report cases to FinCEN by sending a file from the anti-money laundering alert system to the computer system at FinCEN based on FinCEN's communication protocols.

[0222] Conventionally, rule-based systems are used to detect suspicious activities and each rule may trigger an alert. Many financial institutions have used the rule-based approach, which can trigger numerous alerts. For example, there are over two hundred countries in the world. If a financial institution uses a rule-based approach to monitor the wire transfers to, or from, each country, the financial institution may have over two hundred branches at the country decision node of the decision tree. As another example, there are thousands of different industries. If a financial institution uses the rule-based approach to monitor the wire transfers to, or from, each industry, the financial institution may have thousands of branches at the industry decision node of the decision tree. Country and industry are two of many risk categories that have money laundering risks. Similarly, a wire transfer is one of many types of transactions that have money laundering risk. For example, cash, check, ACH, ATM, credit card, debit card, letter of credit, etc., are other possible types of transactions.

[0223] There are many money laundering risk factors. There are numerous (e.g., millions) of possible combinations of branches to form a path from the root of a decision tree to the leaf nodes of the decision tree. In other words, a rule-based system may use millions of rules to cover the entire scope of money laundering risk to detect suspicious money laundering activities. A rule-based system with a limited number of rules may have an increased number of false negatives (e.g., the system has missed true money laundering cases) and many false positives (e.g., the leaf nodes of the decision tree have an increased number of impurities and cannot achieve the goal of classification). Due to the number of false negatives and false positives when a rule-based approach is used, financial institutions hire investigators to review a large number of alerts. It is difficult for financial institutions to mitigate all false negatives with a rule-based system.

[0224] In the field of machine learning, conventional systems consider an accuracy of 70% as satisfactory. It is difficult, if not impossible, to train a machine learning model to have high accuracy, such as 100% accuracy. Unfortunately, although the 70% accuracy may be good for some purposes, this 70% target cannot satisfy the regulatory standards, such as those set by the U.S. government. As discussed, a financial institution may receive a severe regu-

latory penalty if it does not detect certain activity, such as money laundering. As such, a financial institution will not use an alert system with 70% accuracy. Therefore, conventional machine learning models are not satisfactory for an intelligent anti-money laundering alert system.

[0225] According to aspects of the present disclosure, the intelligent anti-money laundering alert system uses a risk-scoring approach. Each risk factor or a degree of a risk factor may be similar to a branch in a rule-based system. As such, the risk scoring process for producing a total risk score from many risk factors, as described in the present disclosure, may consolidate the information from many rules into the total risk score. For example, if a total risk score is generated from 10,000 risk factors, a user only needs to pay attention to those alerts that have the total risk score over a threshold without the need to evaluate each of the 10,000 risk factors. If a rule-based approach is used, each risk factor may have two possible outcomes, matched or not-matched. The total number of possible combinations of outcomes for 10,000 risk factors is two (2) to the power 10,000 (e.g., $2^{10,000}$). Therefore, an evaluation based on the total risk score has effectively replaced the need to evaluate each of the two (2) to the power 10,000 (e.g., $2^{10,000}$) possible outcomes. Because these $2^{10,000}$ outcomes could potentially generate $2^{10,000}$ different types of alerts, the intelligent anti-money laundering alert system can avoid at least $2^{10,000}$ alerts. Therefore, the intelligent anti-money laundering alert system is an improvement in view of the conventional rule-based system.

[0226] Although one total risk score can replace many rules, it may not replace all rules. For example, if a person frequently deposits a certain amount of cash (e.g., \$9,900), which is slightly below the CTR reporting threshold of \$10,000, the financial institution is expected to report this person to Financial Crimes Enforcement Network (FinCEN) as a structuring case. It is difficult to detect a structuring case accurately based on the total risk scores. Therefore, an alert system based on the risk score based technology may include some rules in addition to the criteria based on the risk scores.

[0227] In one aspect of the present disclosure, the intelligent anti-money laundering alert system uses risk score based scenarios to replace rules. In one example, the intelligent anti-money laundering alert system may use approximately twenty to thirty scenarios. The scenarios may include both risk score based scenarios and non-risk score based scenarios.

[0228] In addition to, or alternate from, the scenarios, other conditions may be used to generate an alert. For example, a computer system, such as a machine learning network, may be trained to generate a model. After training, the discriminant used by the model may be converted into an if-then conditional format to trigger alerts.

[0229] For the purposes of this disclosure, a scenario may be defined as a condition or a set of conditions that can trigger an alert or can be used to classify a subject into a category for a particular purpose. For example, a customer with a total risk score within a certain range may not trigger an alert. Still, in this example, the total risk score can classify the customer into a specific risk category, such as high risk, medium risk, or low risk. As another example, a customer that was previously a suspect in a suspicious activity report (SAR) may not trigger an alert. In this example, the customer may be classified into a specific category, such as

prior SAR suspect or another similar category. As another example, a customer matched with the OFAC list, the 314(a) list, the Political Exposed Person list, and/or other lists may be classified in one or more categories.

[0230] A scenario may consist of a rule, a set of rules, a criterion, or a set of criteria, based on rules, facts, behavior patterns, risk scores, risk dimensions, total risk scores, special categories, mathematic models, and/or machine learning models. The scenario may trigger an alert by using a rule-based method, a behavior-based method, a risk-based method, a model-based method, and/or a machine learning-based method (e.g., an artificial intelligence-based method). An intelligent alert system may include one or more scenarios.

[0231] As discussed, an alert may be triggered by a scenario. The scenario may be flagged when one or more conditions are satisfied. A potential case that has triggered an alert may be referred to as a positive. The potential case may include one or more alerts. Therefore, the cause for the potential case may be one or more scenarios. The potential case, or the positive, may be investigated. A true positive may refer to a potential case (e.g., positive) that is a true case. If the investigation indicates that the potential case is not a true case, the potential case may be referred to as a false positive. Consequently, the false positive may be dismissed and the associated alerts may be dismissed as false alerts. The true positive may be reported to an authority, such as FinCEN or law enforcement.

[0232] In one configuration, a posterior probability may be estimated via a Bayesian principle. The product of the posterior probability and evidence is a prior probability multiplied by class likelihood. Using the application of reporting suspicious money laundering activities to FinCEN as an example, the Bayesian equation is $p(S/c)p(c)=p(c/S)p(S)$. The evidence $p(c)$ is the probability of potential cases triggered by the cause c out of all potential cases. The class likelihood $p(S)$ is the probability of the true positives S (e.g., true SAR cases) out of all potential cases. The prior probability $p(c/S)$ is the probability of true positives triggered by the cause c out of all true positives. As a result, the posterior probability $p(S/c)$ may be determined as follows: $p(S/c)=p(c/S)p(S)/p(c)$. The posterior probability $P(S/c)$ is also the conditional probability that a potential case triggered by the cause c is a true positive. That is, although the conditional probability $P(S/c)$ is derived from the historical data, it is the best estimate for the future probability that a potential case triggered by the cause c will become a true positive. As such, the posterior probability may also be referred to as the conditional probability for the future, or future conditional probability.

[0233] Many risk factors (e.g., thousands of risk factors) may affect money laundering risk. In one configuration, the number of scenarios used by the intelligent money laundering alert system is not a large number when the risk score based scenarios are used as a portion of the scenarios. As an example, the intelligent money laundering alert system may use thirty scenarios. A potential case can be triggered by one or more of the scenarios. In this example, a vector with thirty elements may represent the possible causes of the potential case. Consequently, in this example, there are 2^{30} different possible combinations of causes. Each triggered scenario is identified by a flag. For example, a cause vector may be initialized to have a value of "0" for each element. If a

scenario is triggered, the value of an element corresponding to the scenario may change from “0” to another value, such as “1.”

[0234] For example, if a potential case is triggered by a first scenario and a third scenario, the vector, x , may include “1” at the first and third positions, and “0” at all the other positions. That is, the vector can be represented as $x=(1,0,1,0,0, \dots, 0)$. As another example, if a potential case is triggered by a third scenario and a fourth scenario, the third and fourth positions of the vector may include a value of “1,” and all other positions may include a value of “0.” In this example, a vector, x , may be represented as $x=(0,0,1,1,0,0, \dots, 0)$. In the present disclosure, a vector that includes scenarios (e.g., causes) for triggering alerts for a potential case may be referred to as a cause vector.

[0235] A scenario may include one or more conditions for classifying a subject into one or more categories; however, the scenario itself may not trigger a potential case. A potential case may be triggered by multiple scenarios in an associated cause vector. For example, if a scenario intends to classify a subject into the prior SAR suspect category, such a scenario may not trigger a money laundering alert by itself. However, if a customer is a prior SAR suspect and has triggered another scenario (e.g., sending over \$10 million to a higher-risk country), a potential case may be triggered. Still, the cause vector may have two scenarios, one for the money transfer transaction and another for prior SAR suspect. It is a good idea to include various special categories (e.g., prior SAR suspect) into the cause vector as these special categories may improve the accuracy of suspicious activity detection.

[0236] A potential case with multiple triggered scenarios in the cause vector may be more likely to become a true positive. For example, if a customer receives \$250,000 from a wire transfer, one scenario in the cause vector may be flagged (e.g., triggered). This cause vector, with one flagged scenario, may register as a potential case, which may or may not be a true money laundering case. Similarly, if a customer withdraws \$250,000, another scenario in the cause vector may be flagged. Still, this potential case may or may not be a true money laundering case.

[0237] However, if a customer receives \$250,000 from a wire transfer and then withdraws \$250,000 cash from the account, two different scenarios may be flagged in the cause vector. The cause vector, with two flagged scenarios, may register as a potential case, which is more likely to be a true money laundering case because the combined activities described by these two different scenarios match a common money laundering behavior pattern. Therefore, rather than calculating the conditional probability based on one flagged scenario, it is desirable to calculate the conditional probability of a potential case based on a cause vector with multiple flagged scenarios.

[0238] If a cause vector has thirty scenarios, because each scenario has two possibilities (e.g., triggered and not triggered), thirty scenarios may have up to 2^{30} possible combinations. However, because no case will be triggered if none of the scenarios are triggered, the total possible combination for triggering a case is $(2^{30}-1)$. Each combination may have a unique conditional probability to trigger a potential case. It may be impractical to calculate these conditional probability values because 2^{30} is a very large number. In practice, a potential case averages five or fewer simultaneously triggered scenarios. Thus, the actual total

number of meaningful combinations of scenarios that can trigger a potential case is a much smaller number and can be managed via a computational device associated with the intelligent alert system. For example, if the maximum number of possible scenarios in one potential case is five, the total number of possible potential cases that can be triggered by these thirty scenarios is $C(30,1)+C(30,2)+C(30,3)+C(30,4)+C(30,5)$, where $C(m, n)$ is the possible number of different choices to select n objects out of m subjects. For example, $C(30, 1)$ is 30 because there are 30 possible choices to select 1 out of 30 objects. $C(30, 2)$ is 435. $C(30,3)$ is 4,060. $C(30,4)$ is 27,405. $C(30,5)$ is 142,506. The total number of possible cause vectors is 174,436. These cause vectors and their associated conditional probability values can be managed via a computational device and a database associated with the intelligent alert system.

[0239] An investigator may use the intelligent alert system to investigate a potential case triggered by a cause vector. The cause vector may include multiple flagged scenarios. The potential case may be a false positive or true positive. A true positive refers to a potential case that is a true case. A false positive refers to a potential case that is not a true case. If it is a false positive, all the alerts of the potential case are dismissed as false alerts. If it is a true positive, the potential case becomes a true case that may be reported to an authority, such as FinCEN.

[0240] In general, the investigation of one potential case is time consuming. In the U.S., it is common for a large financial institution to hire hundreds of investigators. Each investigator is tasked with investigating whether the potential cases, triggered by various anti-money laundering systems, are true money laundering cases. If there is a true money laundering case, the financial institution is required by U.S. law to report the money laundering case to FinCEN within thirty days. However, as discussed above, whether a potential case is a true money laundering case is a subjective opinion of the investigator.

[0241] If the investigator reports a false positive as a true money laundering case, there is no penalty because the financial institution is protected by the safe harbor rule. In general, because there is a huge regulatory penalty for failing to report a true money laundering case to FinCEN, it is desirable to report a potential case to FinCEN instead of dismissing the potential case. Therefore, it is a general practice for investigators to treat a potential case as a true positive as long as there is reasonable doubt. Current U.S. law does not require the investigator to prove that a potential case is a true case. That is, an investigator is inclined to report a potential case if there is a good chance that it is a true case. This also means that probability plays a role in this decision-making process.

[0242] A user’s decision making may be improved based on knowledge of a conditional probability, $p(S/x)$, of a potential case becoming a true SAR case based on the cause vector x . For example, if the conditional probability is greater than a threshold, the user may want to report the case to FinCEN without spending time on investigation. In one configuration, the intelligent alert system automatically reports a case to an appropriate entity (e.g., FinCEN) when the conditional probability for the case is greater than a threshold. The threshold may be set by a software module, a person who designs or tunes the system, and/or a user of the system. Alternatively, the threshold can be set by the intelligent alert system, which learns the preference of the

user by evaluating the past behavior of the user. For example, if a user often files a SAR when the conditional probability of the cause vector is greater than a value Z , the system can use the value Z as the threshold to automatically file a SAR for the user in the future. In one configuration, the system stores potential cases in a database to determine the conditional probability. For each potential case, the system also stores the associated cause vector. The system may also store the investigation results, such as whether the potential case triggered by the cause vector has been accepted by the investigator as a true positive or has been dismissed by the investigator as a false positive.

[0243] As users continue to use the intelligent alert system, the system accumulates historical data in the database. In one aspect of the present disclosure, for any given period of time, the system may determine, from the database, how many potential cases were triggered by a cause vector, x , and how many of the potential cases triggered by the cause vector, x , have become true positives (e.g., SAR cases reported to FinCEN). The ratio of the number of true positives triggered by the cause vector over the number of potential cases triggered by the cause vector is the conditional probability $p(S/x)$. The conditional probability may also be referred to as a posterior probability. The posterior probability indicates the probability of a future potential case triggered by the cause vector becoming a true case reported to FinCEN. In general, the conditional probability of a potential case is equivalent to the conditional probability of the cause vector that triggered the potential case.

[0244] In one aspect of the present disclosure, the intelligent alert system calculates and displays the conditional probability of each potential case based on its cause vector. The conditional probability indicates the probability of the potential case triggered by the cause vector becoming a true positive reported to FinCEN. In another aspect of the present disclosure, the intelligent alert system accepts a potential case as a true positive and reports it to FinCEN in response to the conditional probability of the cause vector being higher than a predefined value. This predefined value is also referred to as a true positive acceptance threshold.

[0245] The intelligent alert system may also dismiss a potential case as a false positive in response to the conditional probability of the cause vector being less than a false positive rejection threshold. The false positive rejection threshold and true positive acceptance threshold may be set by a software module, a person who designs or tunes the system, and/or user of the system. Alternatively, these thresholds can be set by the intelligent alert system, which learns the preference of the user by evaluating the past behavior of the user. For potential cases that have not been accepted as true positives and have not been dismissed as false positives, the investigators may manually review the potential cases and decide whether each of the potential cases is a false positive or true positive.

[0246] The data for determining the conditional probability may be obtained during a time period. For example, the time period can be the past 12 months, the past three years, or any period. In one configuration, the conditional probability is determined from a rolling time period, which continues to move forward. For example, if the environment (e.g., business policies, customer demographics, products, services, etc.) has changed, the old probability values may not be accurate anymore after the change. Moreover, if the financial institution modifies a scenario, the old probability

values may be affected. Therefore, a rolling time period (e.g., the past three years) provides the intelligent alert system the ability to continue adjusting itself to generate the most current and accurate probability values.

[0247] Many computer systems run data processing batch by batch (e.g., one batch per month). Instead of a time period, a number of batches may also be used to define the amount of historical data that is used for the probability calculations. For example, instead of a rolling time period of the past three years, the computer system can use a rolling time period of the past 36 batches if the computer system runs one batch per month.

[0248] In one configuration, the intelligent alert system purposely leaves some potential cases for the investigators to handle. The intelligent alert system may use the results of these cases to train the system, i.e., adjust the probability values to better fit the current environment. Therefore, the intelligent alert system is a learning system that improves predictions when more potential cases are evaluated by human investigators.

[0249] The intelligent alert system may generate a flag or display a message for a potential case triggered by a cause vector when the cause vector has not generated a potential case during a specified time period. Under such circumstances, a user may manually investigate the potential case to determine whether it is a false positive or true positive. The outcome of the manual investigation may be used to calculate conditional probability values for the cause vector. The calculated conditional probability values may be used to evaluate future potential cases. This manual investigation process has the equivalent effect of supervised training and improves the accuracy and reliability of the intelligent alert system.

[0250] The intelligent alert system may also display or link to the historical potential cases and/or true positives that were triggered by the cause vector. In addition, users may view additional details (e.g., drill down) for each case. Thus, investigators may use historical data as a reference when determining whether to pursue a potential case.

[0251] The system may also display or link to historical potential cases that were triggered by the same suspect of the current potential case and the decisions regarding those potential cases. Investigators may drill down to the detailed background information and transactional information of the suspect. As a result, an investigator can determine whether the current potential case is a false positive or true positive.

[0252] In some cases, there may be insufficient cause for reporting the current potential case to an authority. However, the current potential case combined with historical potential cases together may have sufficient cause for reporting. Under such circumstances, the true cause for reporting the case consists of the cause vector of the current potential case in addition to the cause vectors of the historical potential cases. The historical potential cases may be referred to as prior potential cases. A combined cause vector may be used for this true cause. The combined cause vector may be a combination of multiple cause vectors of multiple potential cases.

[0253] As an example, a cause vector x_1 of a current case may have "1" at the first and fifth position of the vector, and "0" at all other positions (e.g., $x_1=(1, 0, 0, 0, 1, 0, 0, \dots 0)$). In this example, the cause vector x_2 of a historical potential has "1" at the third and fifth position, and "0" at all other positions (e.g., $x_2=(0, 0, 1, 0, 1, 0, 0, \dots 0)$). The combined

cause vector x_3 (e.g., a combination of x_1 and x_2) has “1” at the first position, the third position, and the fifth position, and “0” at all other positions (e.g., $x_3=(1, 0, 1, 0, 1, 0, 0, \dots, 0)$). Although only one cause vector of one historical potential case is used in the above example, a combined cause vector may consist of multiple cause vectors of multiple historical potential cases.

[0254] In one configuration, an investigator has manually reviewed multiple historical potential cases and the current potential case to determine whether the combined case is a false positive (e.g., not to be reported) or a true positive (e.g., to be reported). The outcome of the manual investigation may be used to calculate a conditional probability value $p(S/cbv)$ (e.g., the posterior probability value) for the combined cause vector cbv . The combined cause vector cbv is a combination of the cause vector of the current potential case and one or more cause vectors of historical potential cases.

[0255] In some cases, it is difficult for the intelligent alert system to know which historical potential cases have been investigated by the investigator. As such, the intelligent alert system may prompt the investigator to choose the historical potential cases, which will be combined with the current case to be reported to the authority.

[0256] Additionally, in some cases, it is difficult for the intelligent alert system to know which scenarios of a combined cause vector or a cause vector caused the investigator to report the potential case. Therefore, the intelligent alert system may prompt the investigator to choose the scenarios that caused the investigator to report the potential case.

[0257] Many reports of suspicious activities request the investigator to provide comments or narratives of the potential cases. To improve processing times, it is desirable for the intelligent alert system to automatically populate comments or narratives of the reported cases. In general, the information for writing the comments or narratives consists of the background information and the transactional information of the suspect. Because this information is stored in the database, the Intelligent Alert System can learn from the user how to write the comments or narratives as explained later in this disclosure.

[0258] In one aspect of the present disclosure, the intelligent alert system prompts the investigator to select the historical potential cases to be combined with the current potential case for reporting. Based on the cause vectors of the selected historical potential cases and the cause vector of the current potential case, the intelligent alert system prepares a comment or narrative. The prepared comment or narrative is provided in the report for the combined case.

[0259] When the intelligent alert system fills out the comment or narrative, the intelligent alert system may also identify the combined cause vector of the reported case. Thus, a conditional probability value $p(S/cbv)$ may be associated with the identified combined cause vector cbv based on results of the human investigation.

[0260] The intelligent alert system may prompt the investigator to select the scenarios of a cause vector, or a combined cause vector, that caused reporting of the potential case. Based on the selected scenarios, the intelligent alert system prepares a comment or narrative to fill out a report for the case. These selected scenarios form a real cause vector of the reported case. The scenarios of the real cause vector of the reported case are identified. The conditional probability values of the real cause vector may be calculated based on results of the human investigation.

[0261] Each person may have a unique writing style (or preference), thus, an investigator initially may not like the comment or narrative generated by the intelligent alert system. If the investigator does not like the generated comment or narrative based on the chosen scenarios and there is no way to modify it, the investigator may not bother to select the scenarios that enable the intelligent alert system to generate the comment or narrative. Under such circumstances, the intelligent alert system may not learn the real reasons why the investigator has decided to report the case to the authority. Consequently, the intelligent alert system may not be able to calculate future conditional probability values of the real cause vector based on the human investigation results.

[0262] Therefore, it is desirable for the intelligent alert system to learn and adapt to the writing style (or preference) of the investigator. In one configuration, the intelligent alert system learns the writing style (or preference) of the investigator and generates the future comment or narrative based on the investigator’s writing style (or preference).

[0263] In one configuration, to learn a person’s writing style (or preference), the intelligent alert system displays a comment or narrative for a first chosen scenario based on a pre-stored default comment or narrative for the first chosen scenario. The pre-stored default comment or narrative comprises two primary portions. The first primary portion consists of facts, such as suspect name, identification information, background of the suspect, relationships of the suspect, location of the event, description of the event, date and time of the event, information related to the event, transactional details, etc. The second primary portion may contain words, phrases, sentences, symbols, etc., that are used to link the facts together. These words, phrases, sentences, symbols, etc., are collectively referred to as the “linking words.”

[0264] The facts may be obtained from the stored data or information associated with the intelligent alert system. The investigators may rarely modify the stored facts. An investigator may modify the linking words based on the investigator’s writing style (or preference). Therefore, the intelligent alert system tracks the facts and linking words for the comments and narratives. The intelligent alert system may also track where the facts are stored in memory (e.g., database) and the relationships among the facts.

[0265] In general, the writing style (or preference) of a person is determined by the linking words and the presentation sequence of the facts (e.g., the format). The writing style (or preference) may not be determined based solely on the choices of the facts themselves because the investigator should include the related facts and avoid changing the facts. In some cases, when the same scenario detects two different cases, the facts may be different. Still, the linking words and the sequence of the presentations of the facts (e.g., the format) may remain the same in the comment or narrative because the same investigator has the same writing style (or preference).

[0266] In one configuration, the intelligent alert system provides an editing capability for the investigator to add, delete, or modify the linking words that link the facts together. The intelligent alert system may provide an editing capability for the investigator to add, delete, or modify the facts in the narrative. The intelligent alert system may provide an editing capability and database searching capability for the investigator to extract additional facts from the database and insert them into the narrative.

[0267] After the investigator has revised the comment or narrative for the first chosen scenario, the investigator may store the revised comment or narrative as the next default comment or narrative. In the future, when the investigator selects the first chosen scenario again for other cases, the revised comment or narrative (e.g., the next default comment or narrative) based on a different set of facts may be displayed for the investigator to edit. It is likely that, after a few revisions, the investigator will be satisfied with the then-current revised version and may not want to edit again. Through this evolutionary revision process, the intelligent alert system learns from the investigator and will generate a comment or narrative that matches the writing style (or preference) of the investigator.

[0268] The intelligent alert system can handle a second chosen scenario based on the same approach described above for the first chosen scenario. The intelligent alert system can handle other chosen scenarios through the same manner. Over time, the intelligent alert system will gradually learn how to write a comment or narrative for each scenario based on the investigator's preference.

[0269] As discussed, based on the learning, the intelligent alert system may automatically generate the comment or narrative for the investigator. Based on aspects of the present disclosure, there will be no need for the investigator to write a comment or narrative. The investigator may select the scenarios, and in response, the intelligent alert system automatically fills out the SAR form and the comment or narrative. The intelligent alert system may then report the case to the proper authority. Currently, an investigator may spend several hours writing a comment or narrative for a SAR case. The intelligent alert system can eliminate a tremendous amount of labor for the investigator.

[0270] In some cases, a person's writing may depend on his/her mood. For example, a person in a good mood may write the narrative in detail. As another example, a person in a bad mood may write a poor or incomplete narrative. Aspects of the present disclosure eliminate the influence of the human writer's mood on the narrative, such that the narrative maintains a consistent standard.

[0271] In an exemplary situation, when the intelligent alert system detects a customer, John Doe, deposits \$9,990 on June 1, and \$9,995 on June 2, into an account at Bank ABC, an alert may be generated with a default narrative as follows: "Mr. John Doe deposited \$9,990 on June 1, and \$9,995 on June 2, into Bank ABC." In the short narrative of the present example, the underlined words are the facts and the remaining words are the linking words.

[0272] In one example, an investigator may change the narrative as follows: "Mr. John Doe deposited \$9,990 on June 1, and \$9,995 on June 2, into Bank ABC. According to the Bank Secrecy Act, we are reporting this case as a suspicious activity because this is a typical cash structuring pattern." In the above narrative, the underlined words are the facts and the remaining words are the linking words. When the investigator saves the SAR form on John Doe, the intelligent alert system stores the revised narrative as the default narrative.

[0273] At a later time, the intelligent alert system may detect a customer, Jack Daniel, who deposits \$9,999 on July 1, and \$9,999 on July 2, into an account at Bank ABC. In response, the intelligent alert system may generate a SAR case with the default narrative as follows: "Mr. Jack Daniel deposited \$9,999 on July 1, and \$9,999 on July 2, into Bank

ABC. According to the Bank Secrecy Act, we are reporting this case as a suspicious activity because this is a typical cash structuring pattern."

[0274] In one example, the investigator may change the narrative to the following narrative: "According to the Bank Secrecy Act, a financial institution needs to report a cash structuring activity through Suspicious Activity Report (SAR). We have identified that Mr. Jack Daniel deposited \$9,999 on July 1, and \$9,999 on July 2, into Bank ABC. This is a typical cash structuring activity to avoid the filing of the Currency Transaction Report (CTR). Therefore, we are reporting this case as a suspicious structuring activity case through a SAR." When the investigator saves the SAR form on Jack Daniel, the intelligent alert system stores the revised narrative as the default narrative.

[0275] At a later time period, the intelligent alert system detects a customer, Jim Beam, who deposits \$9,980 on August 3, and \$9,985 on August 4, into an account at Bank ABC. In response, the intelligent alert system may generate a SAR case with the default narrative as follows: "According to the Bank Secrecy Act, a financial institution needs to report a cash structuring activity through Suspicious Activity Report (SAR). We have identified that Mr. Jim Beam deposited \$9,980 on August 3, and \$9,985 on August 4, into Bank ABC. This is a typical cash structuring activity to avoid the filing of the Currency Transaction Report (CTR). Therefore, we are reporting this case as a suspicious structuring activity case through a SAR."

[0276] An investigator may see the above narrative and may want to add a few words as follows: "According to the Bank Secrecy Act, a financial institution needs to report a cash structuring activity through Suspicious Activity Report (SAR). We have identified that Mr. Jim Beam deposited \$9,980 on August 3, and \$9,985 on August 4, into Bank ABC. This is a typical cash structuring activity to avoid the filing of the Currency Transaction Report (CTR). Therefore, we are reporting this case as a suspicious structuring activity case through a SAR. Mr. Jim Beam opened the bank account on Mar. 1, 2019, and the average account balance during the past three (3) months is \$123,197." In this case review process, the investigator has included additional facts that are extracted from the database of the intelligent alert system. These additional facts are underlined in the following sentence: "Mr. Jim Beam opened the bank account on Mar. 1, 2019, and the average account balance during the past three (3) months is \$123,197." When the investigator saves the SAR form on Jim Beam, the intelligent alert system stores the revised narrative as the default narrative.

[0277] At yet another later time period, the intelligent alert system detects a customer, Remy Martin, who deposits \$9,998 on September 5, and \$9,998 on September 6, into an account at Bank ABC. In response, the intelligent alert system may generate a SAR case with the default narrative as follows: "According to the Bank Secrecy Act, a financial institution needs to report a cash structuring activity through Suspicious Activity Report (SAR). We have identified that Mr. Remy Martin deposited \$9,998 on September 5, and \$9,998 on September 6, into Bank ABC. This is a typical cash structuring activity to avoid the filing of the Currency Transaction Report (CTR). Therefore, we are reporting this case as a suspicious structuring activity case through a SAR. Mr. Remy Martin opened the bank account on Feb. 15, 2019, and the average account balance during the past three (3) months is \$83,225."

[0278] The investigator may see the above narrative and decide there is no need to make any changes. Until the investigator makes changes in the future, cases detected by the same scenario will use the following comment or narrative: “According to the Bank Secrecy Act, a financial institution needs to report a cash structuring activity through Suspicious Activity Report (SAR). We have identified that (Suspect Name) deposited (First Cash Transactional Amount) on (First Deposit date), and (Second Cash Transactional Amount) on (Second Deposit Date), into (Bank Name). This is a typical cash structuring activity to avoid the filing of the Currency Transaction Report (CTR). Therefore, we are reporting this case as a suspicious structuring activity case through a SAR. (Suspect Name) opened the bank account on (Account Opening Date), and the average account balance during the past three (3) months is Average Account Balance.” The underlined words will be extracted from the database of the intelligent alert system. The remaining words of the narrative are the preferred linking words used by the investigator that the intelligent alert system has learned from the investigator based on his writing of narrative for the past cases detected by the same scenario.

[0279] In the above example, the set of facts consists of Suspect Name, First Cash Transactional Amount, First Deposit Date, Second Cash Transactional Amount, Second Deposit Date, Bank Name, Account Opening Date, and Average Account Balance. These different pieces of facts can be extracted from a storage location, such as a database.

[0280] Moreover, John Doe, Jack Daniel, Jim Beam, and Remy Martin are the same type of fact under the field name “Suspect Name.” Each suspect name may be defined as a fact corresponding to other suspect names. For example, Remy Martin may be a corresponding piece of fact of Jim Beam. Similarly, a set of corresponding pieces of facts may be defined under the following fields: First Cash Transactional Amount, First Deposit Date, Second Cash Transactional Amount, Second Deposit Date, Bank Name, Account Opening Date, and Average Account Balance.

[0281] When the intelligent alert system shows the default narrative based on a new set of facts of a new suspect, the intelligent alert system replaces each old fact of the old suspect with the new corresponding fact of the new suspect. In the above example, the old suspect name Jim Beam is replaced with the new suspect name Remy Martin; \$9,980 is replaced with \$9,998; August 3 is replaced with September 5; \$9,985 is replaced with \$9,998; August 4 is replaced with September 6; Bank ABC is replaced with Bank ABC; Mar. 1, 2019 is replaced with Feb. 15, 2019; and \$123,197 is replaced with \$83,225. The linking words are unchanged.

[0282] If the investigator has used the same default narrative for a predefined number of times without revising it, the default narrative has matched the writing style (or preference) of the investigator. Under such circumstances, the intelligent alert system may skip or recommend the investigator to skip the narrative reviewing process.

[0283] In one configuration, in addition to providing one comment or narrative for each scenario, the intelligent alert system provides an introduction section for each case. Additionally, or alternatively, the intelligent alert system may provide a conclusion section for each case. The introduction section is placed at the beginning of the overall narrative and the conclusion section is placed at the end of the overall narrative. For example, if a case has three scenarios chosen by the investigator, the overall comment or

narrative will have one introduction section, three comment or narrative sections to match three chosen scenarios, and one conclusion section.

[0284] In one application of the present disclosure, the introduction section and the conclusion section can also be modified and saved by the investigator. Similarly, the intelligent alert system will learn to construct the investigator’s preferred introduction section and conclusion sections. This general format to include an introduction section and a conclusion section provides additional flexibility for the investigator to write a more comprehensive and universal narrative.

[0285] In one configuration, if a case contains multiple suspects, each suspect is detected by a set of scenarios. The overall comment or narrative for the case may contain the introduction section, a relationship section that describes the relationships of these suspects, a single set of comment (or narrative) sections for each scenario, and the conclusion section.

[0286] Updating the linking words and the relative positions of facts in the default narrative based on different sets of facts can simplify the SAR case review and filing process. For example, when the intelligent alert system detects an alert on a suspect, the intelligent alert system sends the current matched scenarios and all the scenarios that match the historical alerts on the suspect to the computer system of the investigator. The investigator selects a scenario constituting the reason to file a SAR and sends the selected scenario back to the intelligent alert system. The intelligent alert system searches the database to identify the default narrative for the selected scenario and sends the default narrative based on the facts of the suspect back to the computer system of the investigator. The investigator reviews the narrative and may make changes, if needed.

[0287] When the investigator saves the revised narrative, the computer system of the investigator sends the revised narrative back to the intelligent alert system. The intelligent alert system stores the revised narrative and sends the SAR form with the revised narrative to the computer system of the BSA Officer. If the BSA Officer approves the SAR form, the intelligent alert system sends the SAR form to the computer system of FinCEN. If the investigator does not see the need to make any change on the default narrative, the intelligent alert system can directly send the SAR with the default narrative to the computer system of the BSA Officer for approval.

[0288] In some cases, the investigator is also the BSA Officer or the BSA Officer permits the investigator to directly file a SAR without any need for approval. In these cases, the investigator may accept the default narrative based on the then-current facts. In response, the intelligent alert system may directly send the SAR with the default narrative based on the current facts to the computer system of FinCEN.

[0289] After the investigator has continuously accepted the default narrative for a scenario based on different sets of facts without any change for a predefined number of times, the intelligent alert system can assume that the default narrative has matched the writing style (or preference) of the investigator for that scenario. Therefore, when a future true-positive case is detected with the same scenario again on a then-current suspect, the intelligent alert system can directly send the SAR with the default narrative based on the then-current facts of the then-current suspect to the com-

puter system of FinCEN. The situation has eliminated the labor associated with the investigator and the BSA Officer.

[0290] The above description for one chosen scenario can also be applied for multiple chosen scenarios. For example, if the investigator has continuously accepted the default narratives for all the chosen scenarios of detected cases based on different sets of facts for a predefined number of times, the intelligent alert system can send the SAR with the default narratives of the multiple chosen scenarios based on the then-current facts of the then-current suspect to the computer system of FinCEN.

[0291] In addition to the SAR filing application, aspects of the present disclosure may be used by a computer system to automatically generate different types of reports based on the preferences of human writers. For example, a hospital may need to generate a report for each patient. A police department may need to generate a report for each incident. A school may need to generate a report for each student. There are many other needs to generate reports. Conventional reports are generated by using tremendous human resources. Aspects of the present disclosure can reduce the human resources used in generating reports.

[0292] Reports may be classified into different types of reports based on different factors, such as reasons, purposes, criteria, scenarios, etc. For example, for a hospital, different types of reports may be used based on the reason a patient checked into the hospital. As an example, the reasons can be heart surgery, baby delivery, etc. A patient may have multiple reasons for checking into a hospital. Additionally, for each main reason, there may be multiple sub-reasons. For example, if a patient checks into the hospital because of a need for a heart surgery, there are many reasons for the need. It is desirable to classify the reasons in detail because each different reason may require a different type of writing style (or preference) to generate a report. As another example, there are many different reasons, purposes, criteria, scenarios, etc., for a police department to generate a report for an incident. In yet another example, there are many different reasons, purposes, criteria, scenarios, etc., for a school to generate a report for each student.

[0293] Reports may be written based on one or more facts. These facts may be stored in a database and consist of the data entered by humans, detected by sensors, collected from different sources, and/or derived from other data. Furthermore, a human will use words, phrases, sentences, symbols, etc., to link the facts together to form a report. For easy reference, the words, phrases, sentences, symbols, etc., to link the facts together are collectively referred to as the "linking words."

[0294] In one configuration, a computer system stores facts in a database. A computer system provides an editing capability for a human writer to create a set of factors, which may comprise reasons, purposes, criteria, scenarios, etc. The computer system may provide an editing capability for a human writer to use a set of facts to create a default narrative for each factor. Additionally, the computer system provides an editing capability for a human writer to write linking words for the default narrative of each factor. The computer system may also store the default narrative of each factor. The default narrative comprises facts and linking words.

[0295] In one configuration, the computer system stores a default narrative for each factor in a database. In this configuration, the default narrative includes linking words, the position of each fact in the narrative, and the storage

location in a database to store each fact. For example, a default narrative can be "(object 1) had a car accident on (object 2)." In this example, object 1 and object 2 are two facts. The computer system stores, in a database, the entire sentence that includes the linking words "had a car accident on" and the positions of object 1 and object 2 in this sentence. In addition, the computer system stores, in the database, the table names and the field names of object 1 and object 2, respectively.

[0296] Data fields with the same definition may be stored in the same database table. For example, all patients' names are stored in the same database table that lists the names of all patients. Therefore, when two different sets of facts are used to write two narratives for two cases, a corresponding pair of facts, which are at the same position inside each respective narrative, are in the same database table. When multiple database tables are used to generate a fact, the database keys to link these multiple database tables can also be stored in the database. As a result, when a default comment or narrative based on an old set of facts is used to generate a new narrative for a new set of facts, the computer system identifies each corresponding pair of facts and replace the old facts with the corresponding new facts.

[0297] For example, object 1 is the "patient name field" stored in the patient table, and object 2 is the "date field" in an event table. In the above example, "Jack Daniel had a car accident on Jan. 20, 2018" and "Jim Beams had a car accident on Feb. 3, 2018" are based on the same narrative format, but contain two different pieces of facts (e.g., patient names and event dates). The linking words of these two scenarios are identical, which are "had a car accident on."

[0298] In one configuration, a computer system lists a set of factors, which may include reasons, purposes, criteria, scenarios, etc. The computer system may permit a human writer to select a factor to display the default narrative based on a new set of facts. The human writer may add, delete, or modify the linking words of the narrative displayed by the computer system.

[0299] In one configuration, the computer system provides database searching and editing capabilities, such that a human writer can add, delete, or modify the facts and change the positions of the facts in a narrative displayed by the computer system. The human writer may store the revised narrative as the new default narrative, which includes the facts, the position of each piece of fact, and the linking words. The computer system stores the database table, key, and field information for obtaining each fact of the new default narrative.

[0300] In one aspect of the present disclosure, a human writer selects a factor to display a new default narrative based on a new set of facts, and the same set of linking words that were stored in the database. The computer system extracts each new piece of the new facts based on where the old corresponding piece of the old facts was stored in the database. The computer system may display each new fact among the linking words in the narrative based on the position of each old corresponding fact in the narrative.

[0301] In one configuration, the computer system provides functionality for a human writer to add, delete, or modify the linking words of the new default narrative displayed by the computer system. The human writer may also add, delete, or modify the facts and change the position of the facts in the new default narrative displayed by the computer system. The

human writer may store the revised new default narrative as the next new default narrative again.

[0302] The above process may be repeated, such that a human writer can continue to revise the default narrative based on a new set of facts and store the revised default narrative as the next new default narrative. As a result of this evolutionary process, the future default narrative may match the preference of the human writer.

[0303] In one aspect of the present disclosure, if the human writer has not changed the narrative for different cases using different sets of facts for a predefined number of instances based on the same factor selected by the human writer, the narrative is considered to have matured for the selected factor. The predefined number can be defined by a person and/or a computer system.

[0304] In one configuration, if the human writer has not changed the linking words displayed by the computer system for different cases using different sets of facts for a predefined number of instances based on the same factors selected by the human writer, the linking words are considered to have matured for the selected factor. The predefined number of instances can be defined by a person and/or a computer system.

[0305] In one configuration, if a narrative has matured for a factor selected by a human writer, the computer system automatically skips or recommends the human writer to skip the narrative reviewing process and uses the current default narrative as the standard narrative format to generate a report for the selected factor. The standard narrative format contains the facts that may be different in each report and the identical set of linking words that match the writing style (or preference) of the human writer.

[0306] In one configuration, if the linking words have matured for a factor selected by a human writer, the computer system automatically skips or recommends the human writer to skip the narrative reviewing process and uses the current default linking words as the standard linking words to generate a report for the selected factor.

[0307] In one configuration, if the human writer has selected multiple factors to write a report, the computer system uses the selected factors to generate one narrative section for each factor and combines multiple narrative sections together based on the multiple selected factors to generate the report.

[0308] An introduction section can be inserted at the front of the report. The introduction section includes facts and/or linking words. The facts and/or linking words may be revised by the human writer through multiple reports to eventually match the writing skill (or preference) of the human writer based on the evolutionary process explained in the present disclosure.

[0309] A linking section may be inserted at the middle of the report. The linking section includes facts and/or linking words, which may be revised by the human writer through multiple reports to eventually match the writing skill (or preference) of the human writer based on the evolutionary process explained in the present disclosure.

[0310] A conclusion section can be inserted at the end of the report. The conclusion section includes facts and/or linking words, which may be revised by the human writer through multiple reports to eventually match the writing skill (or preference) of the human writer based on the evolutionary process explained in the present disclosure.

[0311] As a result of the present disclosure, the computer system learns the writing style (or preference) of each human writer and may automatically generate a variety of reports for each human writer based on his/her writing style (or preference).

[0312] One or more of the examples above are based on the anti-money laundering application in financial institutions. Nevertheless, the present disclosure can also be applied to many other different types of applications for different organizations and different purposes. For example, the intelligent alert system can be used by a government organization to identify any employee who may potentially steal confidential information from the government. The intelligent alert system can be used by a school to identify any student who may potentially drop out of the school. The intelligent alert system can be used by a social network company to identify any member that may potentially conduct illegal activity on the social network. The intelligent alert system can be used by an employer to identify any employee who may potentially quit the job. The intelligent alert system can be used by a marketing company to identify a target for a potential business transaction. The intelligent alert system can also be a mobile application used by an individual to identify a potential stock or commodity for investment purposes. As a public health application, the intelligent alert system can be a mobile app, which monitors the health condition of a person and sends a message if there is a potential health concern. There are countless applications for the intelligent alert system. The following procedure is an example of how to design and develop an intelligent alert system to monitor a group of subjects for any specific goal.

[0313] In one configuration, an intelligent alert system assigns scores to various factors. Additionally, or alternatively, the intelligent alert system assigns scores to each degree of each factor. A degree of a factor is used to distinguish a different level of impact of the factor. For example, sending a wire transfer is a risk factor to be considered for anti-money laundering purposes. However, the dollar amounts of the wire transfers may have different impacts. For example, a wire transfer amount from 0 to \$10,000 may have a low level of money laundering risk while a wire transfer amount from \$250,000 to \$1,000,000 may have a high level of money laundering risk. Factors may be based on data associated with subjects that have positive or negative impacts for achieving the goal. The intelligent alert system assigns a score to each factor. The intelligent alert system may identify the possible degrees of factors in the data associated with the subjects that have positive or negative impacts to achieving the goal. The intelligent alert system assigns a score to each degree of each factor. In one configuration, the intelligent alert system generates a total score for each subject under monitoring by summing all scores of the factors or degrees of factors associated with the subject.

[0314] The intelligent alert system uses a set of scenarios based on different criteria. The criteria may include factors from data associated with the subjects, degrees of factors from data associated with the subjects, and/or scores derived from the data associated with the subjects. Additionally, or alternatively, the criteria may be based on rules derived from a decision tree, special categories associated with the subjects, an if-then conditional format derived from a model trained by the machine learning network, an if-then condi-

tional format derived from a behavior pattern, an if-then conditional format derived from a transactional pattern, factors established by a software module, and/or factors established by a user or designer of the system.

[0315] Through the above methods, the scenarios of the intelligent alert system are established by a variety of methods. These scenarios may trigger alerts to generate potential cases and each potential case may have one scenario or multiple scenarios in its cause vector. The intelligent alert system may list a set of potential cases triggered by one or more scenarios. Investigators may review the potential cases to determine which cases are true positives and which cases are false positives. Additionally, the investigator may review the current potential cases together with the historical potential cases to determine which combinations of cases are true positives or false positives.

[0316] In one configuration, the intelligent alert system enables investigators to review the scenarios of the potential cases to determine which combinations of scenarios generate true positives and which combinations of scenarios generate false positives. The intelligent alert system also provides investigators the ability to review the scenarios of the current potential cases together with the scenarios of the historical potential cases to determine which combinations of scenarios are true positives and which combinations of scenarios are false positives.

[0317] Although a combined cause vector is obtained from the combination of a number of cause vectors, a combined cause vector has the same form of a cause vector. By definition, a combined cause vector is the cause vector of the combined case. Therefore, the conditional probability of a combined cause vector $P(S/cbv)$ and the conditional probability of a cause vector $P(S/x)$ may be calculated via a similar method.

[0318] Furthermore, although a cause vector (or a combined cause vector) may trigger a potential case for investigation, the reason for reporting the case may be based on a subset of scenarios of the cause vector. To maintain accuracy of posterior probability calculation, it is desirable to identify the subset of scenarios that form the real cause vector for the true positive.

[0319] The intelligent alert system provides investigators the ability to review the scenarios of a potential case to identify the real cause vector if the potential case is a true positive. The investigator may review the scenarios of a combined potential case to identify the real cause vector if the combined potential case is a true positive. The intelligent alert system may store the investigation results of each potential case and the associated cause vector (or real cause vector). As explained earlier, once the real cause vector has been identified, the set of scenarios that constitute the real cause vector can be used to generate a set of narratives and the SAR form can be automatically filled out and sent to FinCEN.

[0320] In one configuration, the intelligent alert system stores investigation results of a combined case and the associated combined cause vector (or the real combined cause vector) of the combined case. Each combined cause vector (or the real combined cause vector) may consist of one or more scenarios. The results and other information may be stored in a database or other data structure.

[0321] After the investigators have used the intelligent alert system for a period of time, the intelligent alert system accumulates a large amount of data associated with the

subjects. The data may include historical potential cases, historical investigation results (e.g., true positives or false positives), and the associated cause vectors (or real cause vectors). As a result, the system's accuracy may improve as the system's use increases. That is, the system's accuracy may be improved through the accumulation of data

[0322] For clarity, a cause vector or a real cause vector is generally referred to as a cause vector hereinafter. Moreover, a cause vector generally includes both a cause vector and a combined cause vector hereinafter. Therefore, a cause vector generally refers to a cause vector, a combined cause vector, a real cause vector, and/or a real combined cause vector.

[0323] In one configuration, after an amount of historical data is greater than a threshold, the system calculates the conditional probability for each cause vector. The threshold may be based on a number of true cases, potential cases, data size, and/or other factors. The conditional probability of the cause vector, based on a given time period, is the number of true positives triggered by the cause vector divided by the total number of potential cases triggered by the cause vector.

[0324] In one aspect of the present disclosure, the intelligent alert system dismisses a potential case triggered by a cause vector as a false positive when the conditional probability of the cause vector is lower than a false positive rejection threshold. The false positive rejection threshold may be set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system.

[0325] In some cases, if potential cases triggered by a cause vector always have a low conditional probability, the scenarios of the cause vector may not be properly defined. Under such circumstances, the users adjust the scenarios of the cause vector so that these scenarios will increase probability predictions. The intelligent alert system may prompt the user to make such changes.

[0326] The intelligent alert system may accept a potential case triggered by a cause vector as a true positive in response to the conditional probability of the cause vector being higher than a true positive acceptance threshold. The true positive acceptance threshold may be set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system.

[0327] A vector of multiple elements can be converted into a combination of multiple vectors. For example, vector A has three elements, v_1 , v_2 , and v_3 . In this example, vector A can be the combination of three vectors (e.g., vector B having element v_1 , vector C having element v_2 , and vector D having element v_3). For clarity, vector A is referred to as a parent vector. Vector B, vector C, and vector D may be referred to as child vectors. In the following disclosure, the cause vector will be considered the parent vector.

[0328] The above example assumes that a child vector has only one element. In general, a child vector may have multiple elements. For example, vector A of the above example may have a child vector with elements v_1 and v_2 . Because each element can be included into or excluded from the parent vector to form a child vector, a parent vector with N elements may have a total of 2^N possible combinations, including itself with all N elements and a null vector with no elements. Therefore, a parent vector with N elements may have $2^N - 2$ possible meaningful child vectors. Each element of a cause vector corresponds to a scenario. When the element is one, the corresponding scenario is included. When the element is zero, the corresponding scenario is

excluded. A subset of the scenarios of the parent cause vector may form the scenarios of a child cause vector.

[0329] In general, an increase in the number of scenarios of a cause vector may increase the conditional probability value of the cause vector. For example, if a first cause vector only has scenario A as its vector element and a second cause vector has both scenario A and scenario B as its vector elements, the conditional probability value of the second cause vector should be the same or higher than the conditional probability value of the first cause vector.

[0330] Therefore, a parent cause vector has the same conditional probability value as, or a higher conditional probability value than, any of its child vectors. That is, if a child vector already has a conditional probability value greater than the true positive acceptance threshold, the conditional probability value of the parent cause vector is also greater than the true positive acceptance threshold.

[0331] In one configuration, the intelligent alert system accepts a potential case triggered by a cause vector as a true positive when the conditional probability value of one of its child vectors is equal to or greater than a threshold. The threshold may be set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system.

[0332] A current potential case may be combined with a group of historical potential cases to form a combined cause vector. The intelligent alert system may accept the combined cause vector of the potential cases as a true positive when the conditional probability value of one of the child vectors of the combined cause vector is equal to or greater than a threshold. The threshold may be set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system.

[0333] It may be difficult for the intelligent alert system to try all possible combinations of historical potential cases to determine whether a specific combination of historical potential cases together with the current potential case will satisfy the automatic true positive acceptance criteria. Therefore, in one configuration, the intelligent alert system accepts a combined cause vector as a true positive when the conditional probability value of one of the child vectors of the combined cause vector is equal to or greater than a threshold. The threshold may be set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system.

[0334] In general, all potential cases related to one subject may be related to each other. Additionally, all potential cases related to a group of related subjects may be related to each other. For example, if five students live in the same dormitory, all potential cases related to any one of these five students are related cases. The scope of relationships to define related potential cases may be set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system.

[0335] If the intelligent alert system has been used for an extended period of time, it may not be practical or efficient to use all related potential cases. That is, the number of related potential cases may be too large, thereby decreasing performance. Therefore, it may be desirable to limit the scope of related cases to a period of time. In one configuration, a combined cause vector may be generated from a current potential case and a group of related historical potential cases, which occurred within a predefined time period. The intelligent alert system may accept the combined

cause vector as a true positive when the conditional probability value of a child vector of the combined cause vector is equal to or greater than a threshold. The threshold may be set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system. The predefined period of time is set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system.

[0336] The intelligent alert system provides an opportunity for investigators to investigate cases that are not automatically dismissed as false positives and not automatically accepted as true positives. The intelligent alert system records the investigation results of each potential case and the associated cause vector for the potential case. This information may be used to calculate future conditional probability values of the cause vector.

[0337] Because the intelligent alert system continues to use the investigation results to further adjust future conditional probability values, the intelligent alert system may adjust itself to future environmental changes. The more potential cases the intelligent alert system can handle without human interaction, the fewer potential cases are left for investigators to handle.

[0338] The intelligent alert system may exclude cases that are automatically accepted as true positives or dismissed as false positives from the calculations of the posterior probability values. This approach avoids problems caused by positive feedback. For example, if a potential case triggered by a cause vector x has been automatically accepted as a true positive, the value of the conditional probability $p(S/x)$ may increase if the results of this case are included into the calculations of the posterior probability values of cause vector x . Consequently, the next potential case triggered by the cause vector x may be automatically accepted as a true positive. The automatic acceptance of future potential cases triggered by the cause vector x will continue because the posterior probability value continues to increase. In other words, once a potential case triggered by a cause vector has been automatically accepted as a true positive, all future potential cases triggered by the same cause vector will be automatically accepted as true positives if the accepted cases are included into the calculations of the posterior probability values of the cause vector. This is not desirable because this process of “no return” has deprived the intelligent alert system of the ability to re-adjust itself backward when the environment changes in the future.

[0339] In one configuration, the intelligent alert system does not automatically dismiss a potential case when the conditional probability value of the potential case is lower than the false positive rejection threshold. As a result, an investigator may fine-tune the conditional probability value through this potential case. For reference, this case is referred to as a false positive validation case. The number, percentage, and/or frequency of the occurrence of the false positive validation cases is determined by a software module, a person who designs or tunes the system, and/or a user of the system.

[0340] Additionally, in some cases, the intelligent alert system may not automatically accept a potential case as a true positive when the conditional probability value of the potential case is higher than the true positive acceptance threshold. As a result, an investigator may fine-tune the conditional probability value through this potential case. For clarity, this case is referred to as a true positive validation

case. The number, percentage, and/or frequency of the occurrence of the true positive validation cases is determined by a software module, a person who designs or tunes the system, and/or a user of the system.

[0341] In some cases, certain subjects are handled differently for different reasons. For example, some subjects are placed on a “Do Not Compare List” or a “White List.” A potential case associated with a subject on such a list can be treated as a false positive without any need for an investigation. For example, it may be a politically correct decision to put a political figure on the “Do Not Compare List” of an anti-money laundering system regardless of what has been detected. Similarly, for other purposes, a potential case associated with a subject on another list can be treated as a true positive without any need for an investigation.

[0342] Because these cases are treated differently, they are considered outliers. It is desirable to exclude these outliers from the calculations of the posterior probability values. The intelligent alert system may skip a potential case that is associated with a subject on a “Do Not Compare List” or a “White List.” The skipped case may not be used when calculating posterior probability values of cause vectors.

[0343] In some cases, an alert triggered by a scenario about a subject may turn out to be a false alert because the scenario is not suitable to monitor the subject. For example, a cash-intensive business may naturally have more cash than other types of businesses and a scenario to compare the amount of cash between this business and others may not be meaningful and suitable. Under such circumstances, the investigator can mark the scenario as verified for this subject. This means that the scenario has already been verified by an investigator on this subject and there is no need to take any action if another alert is triggered by this scenario on this subject. Therefore, a potential case triggered by a scenario with a verified status is also considered an outlier.

[0344] In one configuration, the intelligent alert system skips a potential case that is associated with a subject with the verified status on the scenarios that triggered the potential case. The intelligent alert system does not include the skipped case into the calculations of posterior probability values of cause vectors.

[0345] When an investigator dismisses a potential case as a false positive, the intelligent alert system prompts the investigator to determine whether the scenario that triggered the potential case should be marked verified. If this scenario is not marked verified, it may trigger another false positive in the future. Therefore, it is desirable to mark a scenario verified when a potential case triggered by the scenario is determined to be a false positive.

[0346] The number of potential cases used to calculate a conditional probability value may also affect the reliability of the conditional probability value. For example, if only one potential case has been triggered by a cause vector x and the potential case has been accepted by an investigator as true positive, the conditional probability $p(S/x)$ may not be reliable, even though it has a value of 100%. However, if five potential cases have been triggered by a cause vector x and the conditional probability $p(S/x)$ is 100%, this conditional probability may be more reliable in comparison to the previous example.

[0347] The intelligent alert system may automatically dismiss a potential case triggered by a cause vector as a false positive when the conditional probability of the cause vector

is less than a threshold A and the number of potential cases triggered by the cause vector and used to calculate the conditional probability is larger than a threshold B . Each of the thresholds A and B may be set by a software module, a person who designs or tunes the system, and/or a user of the system.

[0348] The intelligent alert system accepts a potential case triggered by a cause vector as a true positive when the conditional probability of the cause vector is higher than the threshold A and the number of potential cases triggered by the cause vector and used to calculate the conditional probability is larger than the threshold B . Each of the thresholds A and B may be set by a software module, a person who designs or tunes the system, and/or a user of the system.

[0349] It may be desirable to use different conditional probability thresholds for subjects in different categories when the intelligent alert system automatically accepts potential cases as true positives or rejects potential cases as false positives based on the conditional probability thresholds. For example, a financial institution may file a SAR on a potential case related to a subject who was a suspect of a past SAR case even though the conditional probability of the current potential case is lower than the true positive acceptance threshold.

[0350] In one configuration, the intelligent alert system uses a different true positive acceptance threshold and false positive rejection threshold for subjects in different categories. The different categories may be defined by a software module, a person who designs or tunes the system, and/or a user of the system. In the anti-money laundering application example, these categories may include a customer who was a suspect of a prior SAR, a customer who had a match with the OFAC list, a customer who had a match with the 314(a) list, a customer who had a match with the political exposed person list, a customer who had a match with other watch lists, higher-risk customer, medium-risk customer, lower-risk customer, higher-risk counter-party, medium-risk counter-party, lower-risk counter-party, higher-risk country, medium-risk country, lower-risk country, higher-risk area, medium-risk area, lower-risk area, higher-transactional amount, medium-transactional amount, lower-transactional amount, etc.

[0351] Because these categories can also be the factors (e.g., risk factors) used for score (e.g., risk score) assignment and calculation purposes, it is desirable to use different true positive acceptance thresholds and false positive rejection thresholds for different factors. In one aspect of the present disclosure, the intelligent alert system permits a user to assign a true positive acceptance threshold and a false positive rejection threshold to each factor.

[0352] In one configuration, the intelligent alert system accepts a potential case as a true positive if the conditional probability of the cause vector is higher than one of the true positive acceptance thresholds of the factors associated with the potential case. The intelligent alert system may reject a potential case as a false positive if the conditional probability of the cause vector is lower than one of the false positive rejection thresholds of the factors associated with the potential case.

[0353] Such an approach may be complicated when many factors are involved. Therefore, it is desirable to only choose some important factors to assign different true positive acceptance thresholds and false positive rejection thresh-

olds. In one configuration, the intelligent alert system permits a user to select a set of factors and assign a true positive acceptance threshold to each selected factor. A user may also select a set of factors and assign a false positive rejection threshold to each selected factor.

[0354] As such, the intelligent alert system may accept a potential case triggered by a cause vector as a true positive if the conditional probability of the cause vector is higher than one of the true positive acceptance thresholds of the selected factors associated with the potential case. Additionally, the intelligent alert system may reject a potential case triggered by a cause vector as a false positive if the conditional probability of the cause vector is lower than one of the false positive rejection thresholds of the selected factors associated with the potential case.

[0355] To increase the accuracy, it is desirable for the total number of potential cases to be greater than a threshold when calculating the conditional probability. The threshold may be a number of cases or a period of time. The threshold may be set by the users as desired.

[0356] In one configuration, the intelligent alert system records the potential case, the investigation results, the associated cause vector, and the date and time when the record is established. The intelligent alert system may calculate the conditional probability of a cause vector x , which is the number of true positives triggered by the cause vector x divided by the total number of potential cases triggered by the cause vector x .

[0357] After calculating the conditional probability value, the intelligent alert system also records into a database additional values, such as: (1) the number of true positives triggered by the cause vector x by that time, (2) the total number of potential cases triggered by the cause vector x by that time, and (3) the date and time of the calculation, which may be referred to as the last calculation time for cause vector x . As a result of storing these additional values, the intelligent alert system does not need to repeat the same calculations to obtain the same values again for the cause vector x .

[0358] The intelligent alert system may update the conditional probability of the cause vector x , which is based on the sum of the number of true positives triggered by the cause vector x (before the last calculation time) and the number of true positives triggered by the cause vector x (after and including the last calculation time) divided by the sum of the total number of potential cases triggered by the cause vector x (before the last calculation time) and the total number of potential cases triggered by the cause vector x (after and including the last calculation time).

[0359] In the above calculations, the number of true positives triggered by the cause vector x (before the last calculation time) plus the number of true positives triggered by the cause vector x (after and including the last calculation time) is the same as the number of true positives triggered by the cause vector x at the time of current calculation. Similarly, the total number of potential cases triggered by the cause vector x (before the last calculation time) plus the total number of potential cases triggered by the cause vector x (after and including the last calculation time) is the same as the total number of potential cases triggered by the cause vector x at the time of current calculation. Therefore, the above calculations will reach the same conditional probability $p(S/x)$, which is the number of true positives triggered by

the cause vector x divided by the total number of potential cases triggered by the cause vector x .

[0360] Both the number of true positives triggered by the cause vector x (before the last calculation time) and the total number of potential cases triggered by the cause vector x (before the last calculation time) may be stored in the database after the last calculation of the conditional probability. Therefore, the intelligent alert system can search the database to find these two values. Therefore, the intelligent alert system calculates two new values based on the potential cases detected after and including the last calculation time. This approach has reduced many calculations, which reduces an amount of data stored in memory.

[0361] In one aspect of the present disclosure, once the calculations of the conditional probability value have been completed, in addition to the potential case, investigation results, and the cause vector x , the intelligent alert system stores additional values, such as: (1) the number of true positives triggered by the cause vector x by that time, (2) the total number of potential cases triggered by the cause vector x by that time, and (3) the date and time of the calculation, which may be referred to as the new last calculation time for cause vector x . As a result, these values will simplify the next round of calculations of the conditional probability for potential cases triggered by cause vector x .

[0362] The above method can be further modified during the software coding process. In one aspect of the present disclosure, the intelligent alert system keeps two counters for a cause vector x , one counter for the number of true positives (NTPX), the other counter for the number of potential cases (NPCX).

[0363] In one aspect of the present disclosure, the intelligent alert system resets both counters, NTPX and NPCX, to zero to start the counting. As one example, a potential case triggered by cause vector x may be manually reviewed by an investigator and determined to be a true positive. In this example, the intelligent alert system adds one to the NTPX counter, because the number of manually reviewed true positives triggered by cause vector x has increased by one. For the current example, the system also adds one to the NPCX counter, because the number of potential cases triggered by cause vector x has increased by one.

[0364] As another example, a potential case triggered by cause vector x is manually reviewed by an investigator and determined to be a false positive. In this example, the intelligent alert system adds zero to the NTPX counter because the number of manually reviewed true positives triggered by cause vector x has not increased, and adds one to the NPCX counter because the number of potential cases triggered by cause vector x has increased by one.

[0365] In one configuration, the conditional probability $p(S/x)$ for a new potential case triggered by cause vector x is NTPX divided by NPCX. This method can reduce the complexity in the calculations of the conditional probability $p(S/x)$ and simplify the software coding effort.

[0366] Although a cause vector x is used in the examples, the above methods can be used for any cause vector. The intelligent alert system may have many pairs of counters, each pair for a cause vector. As explained earlier, the total number of pairs is a limited number because only a very small number of scenarios may coexist in the same cause vector to trigger a potential case.

[0367] By using the above methods, the intelligent alert system may reduce an amount of time for calculations.

Moreover, the conditional probability values increase in accuracy when more potential cases are used in the calculations to derive the conditional probability values.

[0368] Because the intelligent alert system continues to learn from human workers, it is only a matter of time before the intelligent alert system automatically detects an alert, makes the decision to file a SAR, fills out the SAR form, writes a narrative, and sends the SAR form to FinCEN. The intelligent alert system will reduce human resources and will handle the SAR compliance matter similarly to a manner in which a human handles SAR compliance matters.

[0369] Although the detection of suspicious activities, investigation of SAR cases, and filing of suspicious activity reports are used as an example, the same set of methods in the present disclosure can be used to handle the detection of currency transactions, investigation of CTR cases, and filing of Currency Transaction Reports (CTRs) to FinCEN.

[0370] Similarly, the same set of methods in the present disclosure can be used to handle the detection of potential OFAC matches, investigation of the potential matches, and reporting true matches to the Office of the Foreign Assets Control (OFAC). Under such circumstances, the relative correlation (RC) value used to measure the degree of match is equivalent to the risk score used to measure the degree of risk. Therefore, instead of using risk score based scenarios, the intelligent alert system can use RC-based scenarios.

[0371] The OFAC list is just one example of many regulatory lists. The same set of methods in the present disclosure can be used for the detection, investigation, and reporting of matches for all types of regulatory lists, such as 314(a) list, denied persons list, political exposed persons list, and any other lists published by government organizations, and/or non-government organizations. People familiar with the regulatory compliance requirements can appreciate that the set of methods in the present disclosure can be used to detect, investigate, and report any subject to comply with any types of regulatory reporting requirements.

[0372] As discussed, the present disclosure describes a set of methods that can be used by the intelligent alert system for any application and purpose. Whenever an application involves alert generation, alert reviews by humans, and follow-up actions by humans in response to the alert review results, the intelligent alert system will gradually learn from the humans, make decisions on behalf of the humans, and implement the follow-up actions for the humans. As a result, the intelligent alert system will reduce the human effort and time, and may replace some or all humans in such applications.

[0373] As contemplated in the described aspects, one of many possible combinations is described below as an example. An intelligent alert system **500** and a computer network **600**, such as a local area network, enable a BSA officer **100**, compliance officer **200**, investigator **300**, and other responsible person **400** to comply with different types of laws and regulations and send SAR cases directly to another computer system **700** at FinCEN as shown in FIG. 1.

[0374] The compliance officer **200** configures and/or adjusts the parameters of the computer system **500** through the computer network **600**. The computer system **500** uses an internal workflow function to send a potential case through the computer network **600** to the investigator **300**. After the investigation, the investigator **300** sends the potential case and her investigation results through the computer

network **600** to the computer system **500**. The computer system **500** uses an internal workflow function to send the potential case and investigation results through the computer network **600** to the BSA officer **100** for approval. After the BSA officer **100** has approved the investigation results, if the potential case is a true positive, the computer system **500** receives the approval from the BSA officer **100** through the computer network **600**. Then, the computer system **500** sends the true positive to the computer system **700** at FinCEN.

[0375] In some small financial institutions, the same person may have multiple job roles. For example, one person can be the BSA officer, the compliance officer, and the investigator. Under such circumstances, the intelligent alert system uses its internal workflow function to assign different jobs to this person based on his different roles at different stages of the workflow.

[0376] After the computer system **500** has gradually learned the experience of the investigator **300**, the computer system **500** will become smarter and will automatically accept a potential case as a true positive if the conditional probability of the potential case is higher than a predefined value. Under such circumstances, the computer system **500** directly sends the true positive to the computer system **700** at FinCEN without any human involvement. The more the computer system **500** has been used by the investigator **300**, the smarter the computer system **500** becomes. It is expected that the computer system **500** will eventually handle almost all the potential cases by itself with minimum human involvement.

[0377] As shown in the flowchart of FIG. 2 in combination with the system diagram of FIG. 1, the computer system **500** is used for the anti-money laundering application. First (block **2001**), the computer system **500** receives customers' background data and transactional data from a financial institution. Then (block **2002**), the computer system **500** assigns a risk score to each risk factor of the data. The compliance officer **200** has the option to adjust the risk scores through the network **600**. Additionally (block **2003**), the computer system **500** assigns a risk score to each degree of risk factor of the data. The compliance officer **200** also has the option to adjust the risk scores through the network **600**. After the risk scores have been assigned and adjusted, the computer system **500** calculates a total risk score for each customer (block **2004**). Moreover (block **2005**), the computer system **500** establishes a set of risk score based detection scenarios. The compliance officer **200** has the option to adjust the scenarios through the network **600**. Furthermore (block **2006**), the computer system **500** establishes a set of non-risk score based detection scenarios. The compliance officer **200** also has the option to adjust the scenarios through the network **600**. After the scenarios have been established and adjusted, the computer system **500** uses the scenarios to detect potential cases (block **2007**). The computer system **500** uses its workflow function to communicate with the investigator **300** and BSA officer **100** through the network **600** based on the following mechanism (block **2008**): The computer system **500** sends the potential cases to the investigator **300** through the network **600**. The investigator **300** investigates the potential cases and sends the investigation results to the computer system **500** through the network **600**. The computer system **500** sends the investigation results to the BSA officer **100**. The BSA officer **100** approves the investigation results and sends the approval to

the computer system 500 through the network 600. After the approval of the investigation results by the BSA officer 100, the computer system 500 stores the potential cases, the investigation results, and the associated cause vectors in a database (block 2009) with a time stamp. Through the network 600, the BSA officer 100 will instruct the computer system 500 to report the true positives to FinCEN. The computer system 500 will send the true positives to the computer system 700 at FinCEN based on FinCEN's communication protocols.

[0378] FIG. 3 illustrates a flowchart for reporting a potential case for money laundering according to aspects of the present disclosure. According to aspects of the present disclosure, the computer system 500 increases an ability to detect true positive cases after accumulating potential cases, investigation results, and cause vectors (as explained above) for a period of time. The computer system 500 uses a set of true positive acceptance thresholds. The compliance officer 200 approves the true positive acceptance thresholds through the network 600.

[0379] As shown in FIG. 3, at block 3001, the computer system 500 detects a potential case triggered by a cause vector x . Then (block 3002), the computer system 500 calculates a conditional probability, $p(S/x)$, for the cause vector x . The computer system 500 compares the conditional probability value of the cause vector x with a set of true positive acceptance thresholds (decision block 3003). If the conditional probability value of the cause vector x is over any of the thresholds (YES branch 3005), the computer system 500 sends the potential case as a true positive to the computer system 700 at FinCEN (block 3011). If the conditional probability value is not over any of the thresholds (NO branch 3004), the computer system 500 sends the potential case to the investigator 300 for manual investigation (block 3006) through the network 600. The investigator 300 sends the investigation results to the computer system 500 through the network 600. The computer system 500 sends the investigation results to the BSA officer 100 through the network 600. The computer system 500 receives the approval of the investigation results from the BSA officer 100 through the network 600. After the approval, the computer system 500 stores the potential case, the investigation results, the time stamp, and the associated cause vector in a database and the stored information will be used for future calculations of conditional probability values of cause vector x (block 3007). Moreover, the computer system 500 determines whether the investigation results indicate that the potential case is a true positive (decision block 3008). If the potential case is a true positive (YES branch 3010), the computer system 500 sends the true positive to the computer system 700 at FinCEN (block 3011). If the potential case is not a true positive (NO branch 3009), the computer system 500 dismisses the potential case as a false positive.

[0380] In block 3003, the computer system 500 compares the conditional probability value of the cause vector x with a set of true positive acceptance thresholds to determine whether the computer 500 can automatically accept the potential case as a true positive. However, if the total number of potential cases triggered by the cause vector x previously is a very small number, the conditional probability value of the cause vector x may not be reliable. Under such circumstances, the computer system 500 may still send the potential case to investigator 300 for manual investigation (block 3006) through the network 600.

[0381] However, as explained earlier, if the conditional probability value of one of the child vectors of the cause vector x is reliable and is higher than any of the true positive acceptance thresholds, the computer system 500 can still send the potential case as a true positive to the computer system 700 at FinCEN (block 3011).

[0382] As shown in the flowchart of FIG. 4 in combination with the system diagram of FIG. 1, the computer system 500 uses a method to calculate the conditional probability, $p(S/x)$, of a cause vector x based on the investigation results of the investigator 300. The computer system 500 uses a counter NTPX to count the number of true positives triggered by a cause vector x . Additionally, the computer system 500 uses a counter NPCX to count the total number of potential cases triggered by the cause vector x . At the initial stage, the computer system 500 sets the values of both counters to zero. Once a potential case is detected, the computer system 500 determines whether the potential case is triggered by the cause vector x (decision block 4002). If the potential case is not triggered by cause vector x (NO branch 4004), this process ends (block 4005) and the computer system 500 moves to the next potential case. If the potential case is triggered by the cause vector x (YES branch 4003), the computer system 500 determines whether the potential case needs to be handled manually (decision block 4006). As explained earlier, sometimes a potential case may have exceeded the true positive acceptance threshold and there is no need to handle the potential case manually. Sometimes, although a potential case has exceeded the true positive acceptance threshold, the potential case may still need a manual process so the investigator 300 can further enhance the accuracy of the conditional probability value. If the potential case is not to be handled manually (NO branch 4008), the computer system 500 will handle the potential case another way (e.g., automatically sending a true positive to the computer system 700 at FinCEN). Consequently, because there is no manual investigation and the conditional probability is not affected by the potential case, this process ends (block 4005) and the computer system 500 moves to the next potential case.

[0383] If the potential case is to be handled manually (YES branch 4007), the computer system 500 determines whether the potential case needs review, e.g., an investigation (decision block 4009). For example, as explained before, if a customer is on a Do Not Compare List, the computer system 500 should skip the potential case and the conditional probability should not be affected by the skipped case, which is an outlier. Therefore, if there is no need to review the potential case (NO branch 4011), this process ends (block 4005) and the computer system 500 moves to the next potential case.

[0384] If the potential case needs review (YES branch 4010), many events may happen as explained earlier. For example, the computer system 500 sends the potential case to the investigator 300 through the network 600. The investigator 300 sends the investigation results to the computer system 500 through the network 600. The computer system 500 sends the investigation results to the BSA officer 100 through the network 600. The computer system 500 receives an approval of the investigation results from the BSA officer 100 through the network 600. Based on the investigation results, the computer system 500 determines whether the potential case is a true positive (decision block 4012).

[0385] If the potential case is not a true positive (NO branch 4014), the computer system 500 adds zero to the NTPX counter and one to the NPCX counter (block 4016), then stores the investigation results and the date and time of the investigation decision together with the potential case and the cause vector x in the database (block 4017).

[0386] If the potential case is a true positive (YES branch 4013), the computer system 500 adds one to the NTPX counter and one to the NPCX counter (block 4015), then stores the investigation results and the date and time of the investigation decision together with the potential case and the cause vector x in database (block 4017). Then, this process ends (block 4005) and the computer system 500 moves to the next potential case.

[0387] A cause vector of a potential case defines the possible causes to report the potential case. Therefore, two potential cases of two different customers may have the same cause vector. Although cause vector x is used in the above explanation, there may be many cause vectors. Each cause vector may have a flowchart similar to FIG. 4. For example, a potential case triggered by a cause vector y can also be handled in the same way as shown in the flowchart of FIG. 4, except that counter NTPY and counter NPCY are used instead of counter NTPX and counter NPCX.

[0388] Because the method described above has also been identified through the NPCX counter the total number of potential cases triggered by the cause vector x and manually investigated by the investigator 300, the computer system 500 can also determine whether the conditional probability value is reliable. For example, if NPCX has a value of 1 or 2, the conditional probability value of the cause vector x , e.g., NTPX/NPCX, may not be reliable. However, if the NPCX value is greater than a threshold, the conditional probability value of the cause vector x becomes very reliable. Therefore, the computer system 500 can set a true positive acceptance threshold for the cause vector x under the condition that the NPCX value is larger than a predefined value. The predefined value can be set by a software module, a person who designs the system, a person who tunes the system, and/or a user of the system.

[0389] As shown in the flowchart of FIG. 5 in combination with the system diagram of FIG. 1, the computer system 500 gradually learns the writing style of the investigator 300 so that the computer system 500 may automatically generate a report for the investigator 300. As explained earlier, each subject (e.g., a customer) may have its own corresponding set of facts. The positions of the facts relative to the linking words are generally referred to as the format. The linking words and the relative position of each piece of the facts (e.g., format) generally define the writing style of a human writer. A human writer may modify the linking words, the positions of the facts, or add or delete pieces of facts to complete a report.

[0390] First, the computer system 500 generates a report based on a default set of linking words and a default format together with a set of facts of a current subject (block 5002). Then (decision block 5003), the computer system 500 determines whether the investigator 300 modified the previous report generated by the computer system 500. If the investigator 300 did not modify the previous report (NO branch 5004), there is no need for the investigator 300 to make additional modifications on the current report, and the current report can be sent to the BSA Officer 100 for approval

(or alternatively, to the computer system 700 at FinCEN if the investigator 300 has the authority to send the report directly to FinCEN).

[0391] If the investigator 300 modified the previous report generated by the computer system 500 (YES branch 5005), the computer system 500 sends the current report to the investigator 300 for review and modification (block 5006). Then, after the investigator 300 has completed the review, the report is sent back to the computer system 500 (block 5007).

[0392] The computer system 500 then determines whether the investigator 300 made any modifications to the current report (decision block 5008). If the investigator 300 did not modify to the current report (NO branch 5009), the existing default set of linking words and the existing default format have successfully matched the writing style of the investigator 300 and can be used for future reports. If the investigator 300 modified the current report (YES branch 5010), the computer system 500 will use the set of linking words that have been modified by the investigator 300 as the default set of linking words and format (e.g., the positions of the facts) for the next subject (block 5011).

[0393] The above process will repeat for the subsequent subjects. The computer system 500 enables the investigator 300 to continue to modify the linking words and the format for the subsequent subjects until the report generated by the computer system 500 has successfully matched the writing style of the investigator 300, without any need for modification by the investigator 300.

[0394] After the computer system 500 has successfully generated a report that matches the writing style of the investigator 300, the computer system 500 may still periodically send its reports to the investigator 300 so that the computer system 500 can adapt if the investigator 300 changes her writing style. This continuous learning process is desirable to allow the computer system 500 to adapt to changing needs. The frequency of sending reports to the investigator 300 for the purposes of adapting the computer system 500 to changes of writing styles can be determined by a designer of the system, a user of the system, an engineer, or a computer algorithm that adjusts itself based on the past behavior of the investigator 300.

[0395] In this disclosure, a threshold, predefined value, or parameter that may be set by a person, such as the designer, the user, etc., can also be set by the intelligent system that learns the preferences of the person by evaluating past behavior of the person.

[0396] In this disclosure, the terminology “network” generally refers to a communication network or networks, which can be wireless or wired, private or public, real time or non-real time, or a combination of them, and includes the well-known Internet.

[0397] In this disclosure, the terminology “computer” or “computer system” generally refers to either one computer or a group of computers, which may work alone or work together to accomplish the purposes of the system.

[0398] In this disclosure, the terminology “processor” generally refers to either one processor or a group of processors, which may work alone or work together to accomplish the purposes of the processor.

[0399] In this disclosure, the term “module” refers to a single component or multiple components which can be

hardware, software, firmware, or a combination thereof, and may work alone or work together to accomplish the purposes of the module.

[0400] In this disclosure, a “bank” or “financial institution” generally refers to a financial service provider, either a bank or a non-bank, where financial services and money services are provided. Some examples of financial institutions are banks, credit unions, insurance companies, insurance agencies, stockbrokers, stock agencies, bond brokers, bond agencies, commodity brokers, commodity agencies, securities companies, mortgage companies, mortgage agencies, securities companies, money services businesses, agencies for money services businesses, agencies for organizations that provide financial services or money services, financial holding companies, trading companies, trading agencies, other financial service providers, other financial agencies, stock exchanges, commodity exchanges, securities exchanges, currency exchanges, virtual currency companies, virtual currency issuers, virtual currency service providers, virtual currency network providers, virtual currency computer providers, virtual currency dealers, virtual currency exchanges, virtual securities exchanges, bond exchanges, other exchanges, funds managers, investment companies, private equity firms, venture capital firms, merchant acquirers, payment processors, payment card issuers, payment card program managers, internet merchants, transactions processors, securities processors, other organizations related to financial services, etc.

[0401] In this disclosure, a “bank account” or “financial account” generally refers to an account associated with a financial institution, either a bank or a non-bank, where financial transactions can be conducted through financial instruments such as cash, virtual currency, virtual instruments, virtual securities, checks, credit cards, debit cards, ATM cards, stored value cards, gift cards, prepaid cards, wires, monetary instruments, letters of credit, notes, securities, commercial papers, commodities, securities, precious metal, electronic fund transfers, automatic clearing house, etc.

[0402] In this disclosure, “financial transactions” generally refer to transactions related to financial activities, including but not limited to payment, fund transfer, money services, securities issuance, securities exchange, currency exchange, commodity exchange, payroll, invoicing, trading, escrow, insurance, underwriting, merger, acquisition, account opening, account closing, account status check, etc.

[0403] In this disclosure, “trading” generally refers to trading activities, both private and public, including but not limited to trading of stock, currency, virtual currency, virtual instruments, virtual securities, commodities, rights, values, securities, derivatives, goods, services, merchandise, etc.

[0404] In this disclosure, “securities” are generally referred to according to the definition in the Securities Act of 1933 and other laws and regulations related to the Securities Act of 1933. For example, securities may generally include note, stock certificate, bond, debenture, check, draft, warrant, traveler’s check, letter of credit, warehouse receipt, negotiable bill of lading, evidence of indebtedness, certificate of interest or participation in any profit-sharing agreement, collateral-trust certificate, preorganization certificate or subscription, transferable share, investment contract, voting-trust certificate; valid or blank motor vehicle title; certificate of interest in property, tangible or intangible; instrument or document or writing evidencing ownership of

goods, wares, and merchandise, or transferring or assigning any right, title, or interest in or to goods, wares, and merchandise; or, in general, any instrument commonly known as a “security”, or any certificate of interest or participation in, temporary or interim certificate for, receipt for, warrant, or right to subscribe to or purchase any of the foregoing.

[0405] In this disclosure, a “consumer” generally refers to a customer, person, subject, payer, payee, beneficiary, user, or client, etc., seeking to perform a transaction with an individual, an organization, a merchant, and/or a financial institution.

[0406] In this document, the terminology “identification document” generally refers to a passport, driver’s license, voter card, benefits card, student identification card, social security card, national identification card, identity card, certificate of legal status, and other official documents and information bearing instruments that identify a designated individual by certain verifiable characteristics, that are issued or certified by a consulate, embassy, government agency, public or private organizations or other governmental authorities, and that are protected against unauthorized copying or alteration by the responsible party or parties. In particular, such “identification documents” can be formed from various materials, including paper, plastic, polycarbonate, PVC, ABS, PET, Teslin, composites, etc. and can embed the identification information in various formats, including printed or embossed on the document (or card), written on a magnetic medium, programmed into an electronic device, stored in a memory, and combinations thereof. The “identification information” may include, but is not necessarily limited to, names, identification numbers, date of birth, signatures, addresses, passwords, phone numbers, email addresses, personal identification numbers, tax identification numbers, national identification numbers, countries that issue the IDs, states that issue the IDs, ID expiration date, photographs, fingerprints, iris scans, physical descriptions, and other biometric information. The embedded information can be read through optical, acoustic, electronic, magnetic, electromagnetic, and other media.

[0407] In this disclosure, “personal identification information” generally refers to name, address, date of birth, personal identification number, user ID, password, tax identification number, type of the identification document used, identity number associated with the identification document, country, state, government organization and/or a private organization issuing the identification document, expiration date of the identification document, phone number, screen name, e-mail address, photographs, fingerprints, iris scans, physical descriptions, biometrical information, and other information that can be used to identify a person.

[0408] In this disclosure, “personal information” includes personal identification information, personal relationships, personal status, personal background, personal interests, and personal financial information including information related to financial instruments, financial accounts and financial activities, and other information that is related to a person.

[0409] In this disclosure, “financial instruments” generally refer to instruments that are used to conduct financial transactions. Examples of financial instruments include cash, virtual currency, virtual securities, virtual instruments, credit cards, debit cards, ATM cards, prepaid cards, stored value cards, gift cards, checks, monetary instruments, wire

transfers, ACH transfers, letters of credit, notes, securities, commercial papers, commodities, precious metal, gold, silver, etc.

[0410] In this disclosure, a “personal communication device” generally refers to a device interface used for personal communication purposes.

[0411] In this disclosure, a “device interface” generally refers to a keyboard, a keypad, a monitor, a display, a terminal, a computer, a control panel, a vehicle dash board, a network interface, a machinery interface, a video interface, an audio interface, an electrical interface, an electronic interface, a magnetic interface, an electromagnetic interface including electromagnetic wave interface, an optical interface, a light interface, an acoustic interface, a video interface, an audio interface, a contactless interface, a mobile phone interface, a smartphone interface, a smartbook interface, a tablet interface, other communication device interface, a Personal Digital Assistant (PDA) interface, a handheld device interface, a portable device interface, a wireless interface, a wired interface, and other interfaces.

[0412] In this document, the terminology “terminal” or “kiosk” generally refers to equipment, including a computer and/or its peripherals, microprocessor and/or its peripherals, ATM terminal, check-cashing kiosk, money services kiosk, merchant checkout stand, cash register, coin exchange machine, parking lot payment kiosk, other payment kiosks, contactless device, wire line phone, mobile phone, smartphone, smartbook, tablet, personal communication device, tablet device, digital assistant, entertainment device, network interface device, router, and/or Personal Digital Assistant (PDA), etc., which interfaces a user with a computer network, so that the user may interact with computer systems and other equipment connected to the computer network.

[0413] The methodologies described herein may be implemented by various means depending upon the application. For example, these methodologies may be implemented in hardware, firmware, software, or any combination thereof. For a hardware implementation, the processing may be implemented within one or more application specific integrated circuits (ASICs), digital signal processors (DSPs), digital signal processing devices (DSPDs), programmable logic devices (PLDs), field programmable gate arrays (FPGAs), processors, controllers, micro-controllers, microprocessors, electronic devices, other electronic units designed to perform the functions described herein, or a combination thereof.

[0414] For a firmware and/or software implementation, the methodologies may be implemented with modules (e.g., procedures, functions, and so on) that perform the functions described herein. Any machine-readable medium tangibly embodying instructions may be used in implementing the methodologies described herein. For example, software codes may be stored in a memory and executed by a processor. Memory may be implemented within the processor or external to the processor. As used herein the term “memory” refers to any type of long term, short term, volatile, nonvolatile, or other memory and is not to be limited to any particular type of memory or number of memories, or type of media upon which memory is stored.

[0415] If implemented in firmware and/or software, the functions may be stored as one or more instructions or code on a computer-readable medium. Examples include computer-readable media encoded with a data structure and

computer-readable media encoded with a computer program. Computer-readable media includes physical computer storage media. A storage medium may be any available medium that can be accessed by a computer. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM, DVD, or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium that can be used to store desired program code in the form of instructions or data structures and that can be accessed by a computer; disk and disc, as used herein, includes compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk and blue-ray disc where disks usually reproduce data magnetically, while discs reproduce data optically with lasers. Combinations of the above should also be included within the scope of computer-readable media.

[0416] In addition to storage on computer readable medium, instructions and/or data may be provided as signals on transmission media included in a communication apparatus. For example, a communication apparatus may include a transceiver having signals indicative of instructions and data. The instructions and data are configured to cause one or more processors to implement the functions outlined in the claims. The communication apparatus may not store all of the instructions and/or data on a computer readable medium.

[0417] The aspects described in this disclosure can be assembled to form a variety of applications based on the need. Those skilled in the art and technology to which this disclosure pertains can appreciate that alterations and changes in the described structure may be practiced without meaningfully departing from the principal, spirit, and scope of this disclosure. Such alterations and changes should not be construed as deviations from the present disclosure.

What is claimed is:

1. A computer-implemented method for generating a report, comprising:

storing, in a database of a first computer system, a first fact associated with a first subject, a second fact associated with a second subject, and a third fact associated with a third subject, the first fact, the second fact, and the third fact having a same field name in the database;

receiving, at the first computer system from a second computer system, a first report of the first subject, the first report comprising the first fact and a first set of linking words generated by a human writer;

transmitting, from the first computer system to a third computer system, the first report of the first subject;

transmitting, from the first computer system to the second computer system, the second fact and the first set of linking words;

receiving, at the first computer system from the second computer system, a second report of the second subject, the second report comprising the second fact and a second set of linking words generated by the human writer;

transmitting, from the first computer system to the third computer system, the second report of the second subject; and

transmitting, from the first computer system to the third computer system, a third report of the third subject when the first set of linking words corresponds to the second set of linking words, the third report comprising the third fact and the second set of linking words.

2. The computer-implemented method of claim 1, in which the third fact is at least based on customer data.

3. The computer-implemented method of claim 2, in which the customer data is associated with at least one of an individual, an organization, or a combination thereof.

4. The computer-implemented method of claim 2, in which the customer data is associated with at least one of an industry category of a customer, a business type of the customer, a geographical area of the customer, a country of an address of the customer, a nature of a business of the customer, a product type of the business, a services type of the business, a structure of the business, a profession of the customer, a nationality of the customer, a historical record, a type of transaction conducted, a balance of an account, funds inflow, funds outflow, a transactional pattern, a number of transactions, an amount of transactions, a transactional volume, a transactional frequency, a transactional derivative, a location of the transaction, a time of the transaction, a country of the transaction, a sender of a money transfer transaction, a location of the sender, a country of the sender, a nature of the sender, a recipient of a money transfer transaction, a location of the recipient, a country of the recipient, a nature of the recipient, a relationship, a social status, a political exposure, a historical transaction, a number of suspicious activity reports (SARs) filed for money laundering and terrorist financing cases, a category of the first financial institution, a business type of the first financial institution, a geographical area of the first financial institution, a country of a head office of the first financial institution, a nature of a business of the first financial institution, an age of a person, a sex of the person, an income level of the person, an appearance of the person, a judgment about the person, a personal condition of the person, a family condition of the person, a family member of the person, a family member's condition of the person, a friend of the person, a friend's condition of the person, a historical record of the person, an industry category of the person, a geographical area of the person, a country of an address of the person, a profession of the person, a job type of an employee, an education level of an employee, an income level of an employee, a length of employment at a current job, a performance review record, an employment history, a duration of each employment in the employment history, a reason for termination of each employment in the employment history, an age of the employee, a sex of the employee, a personal condition of the employee, a family condition of the employee, a family member of the employee, a family member's condition of the employee, a friend's condition of the employee, a historical record of the employee, a type of work performed, a number of transactions performed, an amount of transactions performed, a largest amount of transaction, a number of transactions with a particular counter party, an amount of transactions with a particular counter party, a number of changes of a crucial record, a number of changes of a crucial record associated with a particular

counter party, a geographical area of an employee's home, a geographical area of an employee's office, a country of an address of the employee, a due diligence result of the customer, a length of an account history, a number of name matches with gambling organizations in transactions, or a combination thereof.

5. The computer-implemented method of claim 1, in which the third fact is at least based on transactional data.

6. The computer-implemented method of claim 5, in which the transaction data is associated with at least one of a person, an organization, or a combination thereof.

7. The computer-implemented method of claim 5, in which the transactional data is associated with at least one of cash, check, wire transfer, ATM (Automated Teller Machine), ACH (Automated Clearing House), virtual currency, virtual security, virtual instrument, credit card, debit card, prepaid card, electronic fund transfer, wires, monetary instruments, letters of credit, notes, securities, commercial papers, commodities, precious metal, account opening, account closure, an account application, deposit, withdrawal, cancellation, balance check, inquiry, credit, debit, or a combination thereof.

8. The computer-implemented method of claim 1, in which the third report comprises a Suspicious Activity Report (SAR).

9. The computer-implemented method of claim 1, in which the second computer system comprises a device interface residing at a financial institution.

10. The computer-implemented method of claim 9, in which the financial institution comprises at least one of a bank, credit union, money services business, financial holding company, insurance company, insurance agency, mortgage company, mortgage agency, stockbroker, stock agency, bond broker, bond agency, commodity broker, commodity agency, trading company, trading agency, other financial service provider, other financial agency, stock exchange, commodity exchange, currency exchange, virtual currency company, virtual currency issuer, virtual currency service provider, virtual currency network provider, virtual currency computer provider, virtual currency dealer, virtual currency exchange, virtual securities exchange, bond exchange, other exchange, funds manager, investment company, private equity firm, venture capital firm, virtual currency company, merchant acquirer, payment processor, payment card issuer, payment card program manager, internet merchant, other organization related to financial services, or a combination thereof.

11. The computer-implemented method of claim 1, in which the third computer system comprises a device interface residing at a government organization.

12. The computer-implemented method of claim 11, in which the government organization comprises a Financial Crimes Enforcement Network (FinCEN).

* * * * *