

entity through private blockchain, receiving a second set of transaction data related to a second set of transactions each involving a payor and a payee and a second set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the second set of transactions. The method may further include updating the knowledge graph with the second set of transaction data and the second set of non-transaction data. The method may further include using the updated knowledge graph to identify hopping transactions in which money is transferred from an initial payor to a final payee through at least one intermediary transaction involving different financial institutions and/or actors other than the first payor and the first payee, the actors acting as payors and/or payees in the intermediary transactions. The method may further include for each account holder, calculating a degree of separation between the account holder and an initial payor, final payee, or actor involved in at least one hopping transaction identified in the updated knowledge graph. The method may further include for each account holder, determining a context score based on the account holder's overall risk score and degree of separation.

[0009] The method may further include using the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to calculate the false positive likelihood of alerts generated by the existing anti-money laundering system.

[0010] The method may further include using the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to perform a key driver analysis to determine factors impacting the false positive likelihood of the alerts.

[0011] The method may further include creating a decision tree from the results of the key driver analysis and ~~using the results of the key driver analysis with a decision tree method to create~~creating decision rules and ~~to determine~~determining new threshold values to replace existing threshold values for the existing anti-money laundering rule based on the decision tree.

[0012] The method may further include performing a gap analysis using a what-if simulation with the new threshold values and the decision rules to determine difference between performance of existing anti-money laundering rules and anti-money laundering rules modified with the new threshold values.

[0013] The method may further include using reinforcement learning determine a recommendation for modifying the existing anti-money laundering rule. The steps may be performed in any order.

[0014] In yet another aspect, the disclosure provides a non-transitory computer-readable medium storing software that may comprise instructions executable by one or more computers which, upon such execution, cause the one or more computers to: (1) build a knowledge graph with a first set of transaction data related to a first set of transactions each involving a payor and a payee, a first set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the first set of transactions, and a plurality of historical alerts each corresponding to a transaction of the first set of transactions; (2) use machine learning to determine overall risk scores for account holders based on the first set of transaction data and the first set of non-transaction data, wherein the account holders are payors and/or payees involved in the first set of transactions; (3) share the overall risk scores, the first set of transaction data, and the first set of non-transaction data with at least one entity through private blockchain; (4) from the at least one entity through private blockchain, receive a second set of transaction data related to a second set of transactions each involving a payor and a payee and a second set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the second set of transactions; (5) update the knowledge graph with the second set of transaction data and the second set of non-transaction data; (6) use the updated knowledge graph to identify hopping transactions in which money is transferred from an initial payor to a final payee through at least one intermediary transaction involving different financial institutions and/or actors other than the first payor and the first payee, the actors acting as payors and/or payees in the intermediary transactions; (7) for each account holder, calculate a degree of separation between the account holder and an initial payor, final payee, or actor involved in at least one hopping transaction identified in the updated knowledge graph; (8) for each account holder, determine a context score based on the account holder's overall risk score and degree of separation; (9) use the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to

calculate the false positive likelihood of alerts generated by the existing anti-money laundering system; (10) use the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to perform a key driver analysis to determine factors impacting the false positive likelihood of the alerts; create a decision tree from the results of the key driver analysis; (11) ~~use the results of the key driver analysis with a decision tree method to~~ create decision rules and ~~to~~ determine new threshold values to replace existing threshold values for the existing anti-money laundering rule based on the decision tree; (12) perform a gap analysis using a what-if simulation with the new threshold values and the decision rules to determine difference between performance of existing anti-money laundering rules and anti-money laundering rules modified with the new threshold values; and (13) use reinforcement learning determine a recommendation for modifying the existing anti-money laundering rule. These steps may be performed in any order.

[0015] In yet another aspect, the disclosure provides a system for modifying an existing rules of an existing AML system to reduce false alerts, which comprises one or more computers and one or more storage devices storing instructions that may be operable, when executed by the one or more computers, to cause the one or more computers to: (1) build a knowledge graph with a first set of transaction data related to a first set of transactions each involving a payor and a payee, a first set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the first set of transactions, and a plurality of historical alerts each corresponding to a transaction of the first set of transactions; (2) use machine learning to determine overall risk scores for account holders based on the first set of transaction data and the first set of non-transaction data, wherein the account holders are payors and/or payees involved in the first set of transactions; (3) share the overall risk scores, the first set of transaction data, and the first set of non-transaction data with at least one entity through private blockchain; (4) from the at least one entity through private blockchain, receive a second set of transaction data related to a second set of transactions each involving a payor and a payee and a second set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the second set of transactions; (5) update the knowledge graph with the second set of transaction data and the second set

of non-transaction data; (6) use the updated knowledge graph to identify hopping transactions in which money is transferred from an initial payor to a final payee through at least one intermediary transaction involving different financial institutions and/or actors other than the first payor and the first payee, the actors acting as payors and/or payees in the intermediary transactions; (7) for each account holder, calculate a degree of separation between the account holder and an initial payor, final payee, or actor involved in at least one hopping transaction identified in the updated knowledge graph; (8) for each account holder, determine a context score based on the account holder's overall risk score and degree of separation; (9) use the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to calculate the false positive likelihood of alerts generated by the existing anti-money laundering system; (10) use the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to perform a key driver analysis to determine factors impacting the false positive likelihood of the alerts; create a decision tree from the results of the key driver analysis; (11) ~~use the results of the key driver analysis with a decision tree method to~~ create decision rules and ~~to~~ determine new threshold values to replace existing threshold values for the existing anti-money laundering rule based on the decision tree; (12) perform a gap analysis using a what-if simulation with the new threshold values and the decision rules to determine difference between performance of existing anti-money laundering rules and anti-money laundering rules modified with the new threshold values; and (13) use reinforcement learning determine a recommendation for modifying the existing anti-money laundering rule. These steps may be performed in any order.

[0016] Other systems, methods, features, and advantages of the disclosure will be, or will become, apparent to one of ordinary skill in the art upon examination of the following figures and detailed description. It is intended that all such additional systems, methods, features, and advantages be included within this description and this summary, be within the scope of the disclosure, and be protected by the following claims.

[0017] While various embodiments are described, the description is intended to be exemplary, rather than limiting, and it will be apparent to those of ordinary skill in

CLAIMS:

We claim:

1. A method for modifying an existing anti-money laundering rule of an existing anti-money laundering system to reduce false alerts, the method comprising:

building a knowledge graph with a first set of transaction data related to a first set of transactions each involving a payor and a payee, a first set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the first set of transactions, and a plurality of historical alerts each corresponding to a transaction of the first set of transactions;

using machine learning to determine overall risk scores for account holders based on the first set of transaction data and the first set of non-transaction data, wherein the account holders are payors and/or payees involved in the first set of transactions;

sharing the overall risk scores, the first set of transaction data, and the first set of non-transaction data with at least one entity through private blockchain;

from the at least one entity through private blockchain, receiving a second set of transaction data related to a second set of transactions each involving a payor and a payee and a second set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the second set of transactions;

updating the knowledge graph with the second set of transaction data and the second set of non-transaction data;

using the updated knowledge graph to identify hopping transactions in which money is transferred from an initial payor to a final payee through at least one intermediary transaction involving different financial institutions and/or actors other than the first payor and the first payee, the actors acting as payors and/or payees in the intermediary transactions;

for each account holder, calculating a degree of separation between the account holder and an initial payor, final payee, or actor involved in at least one hopping transaction identified in the updated knowledge graph;

for each account holder, determining a context score based on the account holder's overall risk score and degree of separation;

using the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to calculate the false positive likelihood of alerts generated by the existing anti-money laundering system;

using the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to perform a key driver analysis to determine factors impacting the false positive likelihood of the alerts;

creating a decision tree from the results of the key driver analysis;

~~using the results of the key driver analysis with a decision tree method to create~~creating decision rules and ~~to determine~~determining new threshold values to replace existing threshold values for the existing anti-money laundering rule based on the decision tree;

performing a gap analysis using a what-if simulation with the new threshold values and the decision rules to determine difference between performance of existing anti-money laundering rules and anti-money laundering rules modified with the new threshold values; and

using reinforcement learning determine a recommendation for modifying the existing anti-money laundering rule.

2. The method of claim 1, wherein determining a context score based on the account holder's overall risk score and degree of separation includes using the knowledge graph to create a plurality of bipartite adjacency matrices in which a first set of vertices include payors and a second set of vertices include payees.

3. The method of claim 2, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes building a plurality of bipartite graphs from the plurality of bipartite adjacency matrices.

4. The method of claim 3, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes detecting a plurality of communities within the plurality of bipartite graphs.
5. The method of claim 1, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes finding communities of payors and/or payees having similar behavior patterns within the knowledge graph and determining an overall risk score by using a first machine learning model to determine a first risk score for each of the plurality of communities, wherein the first risk score is calculated for each of the following types of transaction data: transaction location, transaction type, number of transactions, and transaction amount.
6. The method of claim 5, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes using a second machine learning model to determine a second risk score for at least a portion of the plurality of payors and/or payees, wherein the second risk score is calculated for each of the following types of transaction data: transaction location, transaction type, number of transactions, and transaction amount, and using an ensemble model of the first and second machine learning models to determine overall risk scores for the at least a portion of the plurality of payors and/or payees.
7. The method of claim 1, wherein the payor and/or payee is one of an individual, an entity, a bank account.
8. The method of claim 1, further comprising modifying the existing anti-money laundering rule according to the recommendation.

9. A system for modifying an existing anti-money laundering rule of an existing anti-money laundering system to reduce false alerts, comprising:

one or more computers and one or more storage devices storing instructions that are operable, when executed by the one or more computers, to cause the one or more computers to:

build a knowledge graph with a first set of transaction data related to a first set of transactions each involving a payor and a payee, a first set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the first set of transactions, and a plurality of historical alerts each corresponding to a transaction of the first set of transactions;

use machine learning to determine overall risk scores for account holders based on the first set of transaction data and the first set of non-transaction data, wherein the account holders are payors and/or payees involved in the first set of transactions;

share the overall risk scores, the first set of transaction data, and the first set of non-transaction data with at least one entity through private blockchain;

from the at least one entity through private blockchain, receive a second set of transaction data related to a second set of transactions each involving a payor and a payee and a second set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the second set of transactions;

update the knowledge graph with the second set of transaction data and the second set of non-transaction data;

use the updated knowledge graph to identify hopping transactions in which money is transferred from an initial payor to a final payee through at least one intermediary transaction involving different financial institutions and/or actors other than the first payor and the first payee, the actors acting as payors and/or payees in the intermediary transactions;

for each account holder, calculate a degree of separation between the account holder and an initial payor, final payee, or actor involved in at least one hopping transaction identified in the updated knowledge graph;

for each account holder, determine a context score based on the account holder's overall risk score and degree of separation;

use the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to calculate the false positive likelihood of alerts generated by the existing anti-money laundering system;

use the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to perform a key driver analysis to determine factors impacting the false positive likelihood of the alerts;

create a decision tree from the results of the key driver analysis;

~~use the results of the key driver analysis with a decision tree method to~~ create decision rules and ~~to~~ determine new threshold values to replace existing threshold values for the existing anti-money laundering rule based on the decision tree;

perform a gap analysis using a what-if simulation with the new threshold values and the decision rules to determine difference between performance of existing anti-money laundering rules and anti-money laundering rules modified with the new threshold values; and

use reinforcement learning determine a recommendation for modifying the existing anti-money laundering rule.

10. The system of claim 9, wherein determining a context score based on the account holder's overall risk score and degree of separation includes using the knowledge graph to create a plurality of bipartite adjacency matrices in which a first set of vertices include payors and a second set of vertices include payees.

11. The system of claim 10, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes building a plurality of bipartite graphs from the plurality of bipartite adjacency matrices.

12. The system of claim 11, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes detecting a plurality of communities within the plurality of bipartite graphs.

13. The system of claim 9, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes finding communities of payors and/or payees having similar behavior patterns within the knowledge graph and determining an overall risk score by using a first machine learning model to determine a first risk score for each of the plurality of communities, wherein the first risk score is calculated for each of the following types of transaction data: transaction location, transaction type, number of transactions, and transaction amount.

14. The system of claim 13, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes using a second machine learning model to determine a second risk score for at least a portion of the plurality of payors and/or payees, wherein the second risk score is calculated for each of the following types of transaction data: transaction location, transaction type, number of transactions, and transaction amount, and using an ensemble model of the first and second machine learning models to determine overall risk scores for the at least a portion of the plurality of payors and/or payees.

15. The system of claim 9, wherein the payor and/or payee is one of an individual, an entity, a bank account.

16. The system of claim 9, further comprising modifying the existing anti-money laundering rule according to the recommendation.

17. A non-transitory computer-readable medium storing software comprising instructions executable by one or more computers which, upon such execution, cause the one or more computers to:

build a knowledge graph with a first set of transaction data related to a first set of transactions each involving a payor and a payee, a first set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the first set of transactions, and a plurality of historical alerts each corresponding to a transaction of the first set of transactions;

use machine learning to determine overall risk scores for account holders based on the first set of transaction data and the first set of non-transaction data, wherein the account holders are payors and/or payees involved in the first set of transactions;

share the overall risk scores, the first set of transaction data, and the first set of non-transaction data with at least one entity through private blockchain;

from the at least one entity through private blockchain, receive a second set of transaction data related to a second set of transactions each involving a payor and a payee and a second set of non-transaction data related to at least one payor and/or at least one payee involved in at least one of the second set of transactions;

update the knowledge graph with the second set of transaction data and the second set of non-transaction data;

use the updated knowledge graph to identify hopping transactions in which money is transferred from an initial payor to a final payee through at least one intermediary transaction involving different financial institutions and/or actors other than the first payor and the first payee, the actors acting as payors and/or payees in the intermediary transactions;

for each account holder, calculate a degree of separation between the account holder and an initial payor, final payee, or actor involved in at least one hopping transaction identified in the updated knowledge graph;

for each account holder, determine a context score based on the account holder's overall risk score and degree of separation;

use the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to

calculate the false positive likelihood of alerts generated by the existing anti-money laundering system;

use the context scores, historical alert data, one or both of the first and second transaction data, and one or both of the first and second non-transaction data to perform a key driver analysis to determine factors impacting the false positive likelihood of the alerts;

create a decision tree from the results of the key driver analysis;

~~use the results of the key driver analysis with a decision tree method to create~~ decision rules and ~~to~~ determine new threshold values to replace existing threshold values for the existing anti-money laundering rule based on the decision tree;

perform a gap analysis using a what-if simulation with the new threshold values and the decision rules to determine difference between performance of existing anti-money laundering rules and anti-money laundering rules modified with the new threshold values; and

use reinforcement learning determine a recommendation for modifying the existing anti-money laundering rule.

18. The non-transitory computer-readable medium storing software of claim 17, wherein determining a context score based on the account holder's overall risk score and degree of separation includes:

using the knowledge graph to create a plurality of bipartite adjacency matrices in which a first set of vertices include payors and a second set of vertices include payees; and

building a plurality of bipartite graphs from the plurality of bipartite adjacency matrices.

19. The non-transitory computer-readable medium storing software of claim 17, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes finding communities of payors and/or payees having similar behavior patterns within the knowledge graph and determining an overall risk score by using a first machine learning model to determine a first risk score for each of the plurality of communities, wherein the first risk score is calculated for each of the following types of transaction data: transaction location, transaction type, number of transactions, and transaction amount.

20. The non-transitory computer-readable medium storing software of claim 17, wherein determining a context score based on the account holder's overall risk score and degree of separation further includes using a second machine learning model to determine a second risk score for at least a portion of the plurality of payors and/or payees, wherein the second risk score is calculated for each of the following types of transaction data: transaction location, transaction type, number of transactions, and transaction amount, and using an ensemble model of the first and second machine learning models to determine overall risk scores for the at least a portion of the plurality of payors and/or payees.