

CLAIMS

1. A computer-implemented method comprising:
 - obtaining, by a server system, output data generated by a first transaction model deployed at a first end-user device, wherein:
 - the first transaction model is trained to apply a set of evidence factors to predict a likelihood that first transaction data received by the first end-user device, or a first target entity transmitting the first transaction data from the first end-user device, is anomalous with respect to a plurality of peer groups, and
 - each peer group included in the plurality of peer groups is associated with a set of attributes, wherein:
 - (i) the set of attributes include an entity classification,
 - (ii) each peer group is associated with the entity classification, and
 - (iii) each peer group is associated with a level of similarity shared between the set of attributes;
 - evaluating, by the server system, the output data in relation to the plurality of peer groups;
 - classifying, by the server system and based on the set of attributes corresponding to the entity classification of the first target entity, the first target entity or a second target entity as belonging to at least:
 - (i) a first peer group of entities that share a first set of attributes and
 - (ii) a second group of entities that share a second set of attributes;
 - processing, by the server system and using the trained first transaction model, the transaction data for the first target entity or the second target entity in relation to the transaction data for entities included in the first peer group to determine a first indicator for the first target entity or the second target entity, wherein the first indicator represents a first classification of normalcy relative to transaction data of the first peer group with respect to an activity pattern type given the first set of attributes, wherein:
 - (i) a first classification of normalcy corresponds to expected activity associated with the transaction data of the first peer group and

(ii) the first target entity is determined anomalous when an activity of the first target entity deviates beyond a threshold from the expected activity of other entities in the first peer group of entities;

determining, by the server system, based on the first indicator, that the transaction data for the first target entity, having the set of attributes corresponding to the entity classification, and the second target entity is (i) potentially anomalous relative to the first classification of normalcy relative to the transaction data and with respect to the activity pattern type and (ii) not potentially anomalous relative to a second classification of normalcy represented by a second indicator indicating the second classification of normalcy relative to transaction data of a second peer group with respect to the activity pattern type given the second set of attributes, wherein a first classification of normalcy corresponds to expected activity associated with the transaction data of the first peer group;

in response to determining that the transaction data is potentially anomalous relative to the first classification of normalcy and is not potentially anomalous relative to the second classification of normalcy:

adjusting, by the server system, the activity pattern type;

generating, by the server system, a second set of evidence factors based on the adjusted activity pattern type;

re-training, by the server system, the first transaction model using the second set of evidence factors to identify potentially anomalous activity associated with the adjusted activity pattern type;

generating, by the server system and based on the re-training of the first transaction model, a model update specifying an adjustment for a second transaction model deployed at a second end-user device, wherein the second transaction model is initially trained to apply the first set of evidence factors, the model update specifying a change to the second transaction model derived from the re-training of the first transaction model; and

providing, by the server system, the model update for output to the second end-user device, wherein the model update, when received by the second end-user device, causes the second end-user device to adjust the second transaction model according to the change specified by the model update.

2. The method of claim 1, wherein the output data comprises:
 - a set of model parameters associated with the first target entity; and
 - a respective model weight for each model parameter included in the set of model parameters.
3. The method of claim 2, wherein:
 - The first transaction model is trained to identify the potentially anomalous activity associated with the first target entity by applying the set of evidence factors to transaction data associated with the first target entity; and
 - the output data excludes the transaction data associated the first target entity.
4. The method of claim 2, wherein:
 - the second transaction model is trained to identify the potentially anomalous activity associated with the second target entity based on a second set of model parameters; and
 - the adjustment for the second transaction model comprises removing a model parameter included in the set of model parameters associated with the first target entity and the second set of model parameters associated with the second target entity.
5. The method of claim 2, wherein:
 - the second transaction model is trained to identify the potentially anomalous activity associated with the second target entity based on a second set of model parameters; and
 - the adjustment for the second transaction model comprises adding a model parameter included in the set of model parameters associated with the first target entity and not included in the second set of model parameters associated with the second target entity.
6. The method of claim 2, wherein:
 - the second transaction model is trained to identify the potentially anomalous activity associated with the second target entity based on a second set of model parameters; and

the adjustment for the second transaction model comprises adjusting a model weight of a model parameter included in the set of model parameters associated with the first target entity and the second set of model parameters associated with the second target entity.

7. The method of claim 1, wherein:

the output data comprises a peer group of entities that share the first set of attributes with the first target entity and the second target entity; and

the adjustment for the second transaction model comprises a change to a peer group classification for the second target entity.

8. The method of claim 7, further comprising:

receiving, by the server system and from the second end-user device, an indication that the second transaction model was adjusted based on the model update; and

based on receiving the indication that the second transaction model was adjusted based on the model update, providing, by the server system and to the second end-user device, an instruction, that when received by the second end-user device, causes the second end-user device to perform operations comprising;

classifying second target entity to the peer group classification that was changed based on the model update;

obtaining, from one or more data sources, transaction data for the first target entity and transaction data for entities included in the peer group classification that was changed based on the model update;

processing, using the second transaction model, the transaction data for the second target entity in relation to the transaction data for entities included in the peer group classification that was changed based on the model update to determine a prioritization indicator for the second target entity; and

enabling a user to perceive a representation of the prioritization indicator.

9. The method of claim 8, wherein:

the prioritization indicator comprises a score; and

a value of the score represents a number of potentially anomalous transactions included in the transaction data for the second target entity.

10. The method of claim 9, wherein a value of the score represents a probability that set of transactions of the second target entity are determined to be anomalous relative to entities included in the peer group that was adjusted based on the model update and to which the second target entity was classified.

11. The method of claim 1, further comprising:

receiving, by the server system and from the second end-user device, an indication that the second transaction model has been adjusted according to the change specified by the model update;

obtaining, by the server system, second output data generated by second transaction model after being adjusted according to the change specified by the model update;

determining, by the server system, an adjustment for a third transaction model deployed at a third end-user device, wherein the third transaction model is trained to apply the set of evidence factors to identify potentially anomalous activity associated with a third target entity determined to be similar to the second target entity;

generating, by the server system, a model update for the third transaction model, wherein the model update for the third transaction model specifies a change to the third transaction model; and

providing, by the server system, the model update for the third transaction model for output to the third end-user device, wherein the model update, when received by the third end-user device, causes the third end-user device to adjust the third transaction model according to the change specified by the model update for the third transaction model.

12. The method of claim 1, wherein:

evaluating the output data in relation to the plurality of peer groups comprises determining a number of peer groups specified within the output data; and

determining the adjustment for the second transaction model comprises:

determining that the number of peer groups specified within the output data satisfies a first threshold, and

determining to adjust the second transaction model based on determining that the number of peer groups specified within the output data satisfies the first threshold.

13. The method of claim 1, wherein:

evaluating the output data in relation to the plurality of peer groups comprises determining a number of sets of attributes specified within the output data; and determining the adjustment for the second transaction model comprises:

determining that the number of sets of attributes specified within the output data satisfies a second threshold, and

determining to adjust the second transaction model based on determining that the number of peer groups specified within the output data satisfies the second threshold.

14. A system comprising:

one or more computing devices; and

one or more storage devices storing instructions that, when executed by the one or more computing devices, causes the one or more computing devices to perform operations comprising:

obtaining, by the one or more computing devices, output data generated by a first transaction model deployed at a first end-user device, wherein:

the first transaction model is trained to apply a set of evidence factors to predict a likelihood that first transaction data received by the first end-user device, or a first target entity transmitting the first transaction data from the first end-user device, is anomalous with respect to a plurality of peer groups, and

each peer group included in the plurality of peer groups is associated with a set of attributes, wherein:

(i) the set of attributes include an entity classification,

(ii) each peer group is associated with the entity classification,

and

(iii) each peer group is associated with a level of similarity shared between the set of attributes;

evaluating, by the one or more computing devices, the output data in relation to the plurality of peer groups;

classifying, by the one or more computing devices and based on the set of attributes corresponding to the entity classification of the first target entity, the first target entity or a second target entity as belonging to at least (i) a first peer group of entities that share a first set of attributes and (ii) a second group of entities that share a second set of attributes;

processing, by the one or more computing devices and using the trained first transaction model, the transaction data for the first target entity or the second target entity in relation to the transaction data for entities included in the first peer group to determine a first indicator for the first target entity or the second target entity, wherein the first indicator represents a first classification of normalcy relative to transaction data of the first peer group with respect to an activity pattern type given the first set of attributes, wherein:

- (i) a first classification of normalcy corresponds to expected activity associated with the transaction data of the first peer group and
- (ii) the first target entity is determined anomalous when an activity of the first target entity deviates beyond a threshold from the expected activity of other entities in the first peer group of entities;

determining, by the one or more computing devices, based on the first indicator, that the transaction data for the first target entity, having the set of attributes corresponding to the entity classification, and the second target entity is (i) potentially anomalous relative to the first classification of normalcy relative to the transaction data and with respect to the activity pattern type and (ii) not potentially anomalous relative to a second classification of normalcy represented by a second indicator indicating the second classification of normalcy relative to transaction data of a second peer group with respect to the activity pattern type given the second set of attributes, wherein a first classification of normalcy corresponds to expected activity associated with the transaction data of the first peer group;

in response to determining that the transaction data is potentially anomalous relative to the first classification of normalcy and is not potentially anomalous relative to the second classification of normalcy:

adjusting, by the one or more computing devices, the activity pattern type;

generating, by the one or more computing devices, a second set of evidence factors based on the adjusted activity pattern type;

re-training, by the one or more computing devices, the first transaction model using the second set of evidence factors to identify potentially anomalous activity associated with the adjusted activity pattern type;

generating, by the one or more computing devices and based on the re-training of the first transaction model, a model update specifying an adjustment for a second transaction model deployed at a second end-user device, wherein the second transaction model is initially trained to apply the first set of evidence factors, the model update specifying a change to the second transaction model derived from the re-training of the first transaction model; and

providing, by the one or more computing devices, the model update for output to the second end-user device, wherein the model update, when received by the second end-user device, causes the second end-user device to adjust the second transaction model according to the change specified by the model update.

15 The system of claim 14, wherein the output data comprises:

a set of model parameters for identifying the potentially anomalous activity associated with the first target entity; and

a respective model weight for each model parameter included in the set of model parameters.

16. The system of claim 15, wherein:

The first transaction model is trained to identify the potentially anomalous activity associated with the first target entity by applying the set of evidence factors to transaction data associated with the first target entity; and

the output data excludes the transaction data associated the first target entity.

17. At least one non-transitory computer-readable storage device storing instructions that, when executed by one or more processors, causes the one or more processors to perform operations comprising:

obtaining output data generated by a first transaction model deployed at a first end-user device, wherein:

the first transaction model is trained to apply a set of evidence factors to predict a likelihood that first transaction data received by the first end-user device, or a first target entity transmitting the first transaction data from the first end-user device, is anomalous with respect to a plurality of peer groups, and

each peer group included in the plurality of peer groups is associated with a set of attributes, wherein:

- (i) the set of attributes include an entity classification,
- (ii) each peer group is associated with the entity classification, and
- (iii) each peer group is associated with a level of similarity shared between the set of attributes;

evaluating the output data in relation to the plurality of peer groups;

classifying, by the one or more processors and based on the set of attributes corresponding to the entity classification of the first target entity, the first target entity or a second target entity as belonging to at least (i) a first peer group of entities that share a first set of attributes and (ii) a second group of entities that share a second set of attributes;

processing, by the one or more processors and using the trained first transaction model, the transaction data for the first target entity or the second target entity in relation to the transaction data for entities included in the first peer group to determine a first indicator for the first target entity or the second target entity, wherein the first indicator represents a first classification of normalcy relative to transaction data of the first peer group with respect to an activity pattern type given the first set of attributes, wherein:

- (i) a first classification of normalcy corresponds to expected activity associated with the transaction data of the first peer group and

(ii) the first target entity is determined anomalous when an activity of the first target entity deviates beyond a threshold from the expected activity of other entities in the first peer group of entities;

determining, by the one or more processors, based on the first indicator, that the transaction data for the first target entity, having the set of attributes corresponding to the entity classification, and the second target entity is (i) potentially anomalous relative to the first classification of normalcy relative to the transaction data and with respect to the activity pattern type and (ii) not potentially anomalous relative to a second classification of normalcy represented by a second indicator indicating the second classification of normalcy relative to transaction data of a second peer group with respect to the activity pattern type given the second set of attributes, wherein a first classification of normalcy corresponds to expected activity associated with the transaction data of the first peer group;

in response to determining that the transaction data is potentially anomalous relative to the first classification of normalcy and is not potentially anomalous relative to the second classification of normalcy:

adjusting, by the one or more processors, the activity pattern type;

generating, by the one or more processors, a second set of evidence factors based on the adjusted activity pattern type;

re-training, by the one or more processors, the first transaction model using the second set of evidence factors to identify potentially anomalous activity associated with the adjusted activity pattern type;

generating, by the one or more processors and based on the re-training of the first transaction model, a model update specifying an adjustment for a second transaction model deployed at a second end-user device, wherein the second transaction model is initially trained to apply the first set of evidence factors, the model update specifying a change to the second transaction model derived from the re-training of the first transaction model; and

providing, by the one or more processors, the model update for output to the second end-user device, wherein the model update, when received by the second end-user device, causes the second end-user device to adjust the second transaction model according to the change specified by the model update.

18. The non-transitory computer-readable storage device of claim 17, wherein the output data comprises:

- a set of model parameters associated with the first target entity; and
- a respective model weight for each model parameter included in the set of model parameters.

19. The non-transitory computer-readable storage device of claim 18, wherein:

- The first transaction model is trained to identify the potentially anomalous activity associated with the first target entity by applying the set of evidence factors to transaction data associated with the first target entity; and

- the output data excludes the transaction data associated the first target entity.

20. The non-transitory computer-readable storage device of claim 18, wherein:

- the second transaction model is trained to identify the potentially anomalous activity associated with the second target entity based on a second set of model parameters; and

- the adjustment for the second transaction model comprises adding a model parameter included in the set of model parameters associated with the first target entity and not included in the second set of model parameters associated with the second target entity.