

Claims

1. A method of training a supervised machine learning system to detect anomalies within transaction data, the method comprising:

obtaining a training set of data samples, each data sample being derived, at least in part, from transaction data and being associated with one of a set of uniquely identifiable entities, wherein at least a portion of the training set is unlabelled;

assigning a label indicating an absence of an anomaly to unlabelled data samples in the training set;

partitioning the data of the data samples in the training set into two feature sets, a first feature set representing observable features and a second feature set representing context features, the observable features being derived from a function of at least transaction data for a current transaction, the context features being derived from one or more of a function of historical transaction data that excludes the current transaction and retrieved data relating to the uniquely identifiable entity for the current transaction;

generating synthetic data samples by combining features from the two feature sets that respectively relate to two different ones of the set of uniquely identifiable entities;

assigning a label indicating a presence of an anomaly to the synthetic data samples;

augmenting the training set with the synthetic data samples; and

training a supervised machine learning system with the augmented training set and the assigned labels,

wherein the trained supervised machine learning system is configured to output a value indicative of a presence of an anomaly when supplied with a new data sample.

2. The method of claim 1, wherein the observable features are derived from a function of transaction data within a predefined temporal window for the current transaction and wherein the context features are derived from transaction data outside of the predefined temporal window.

3. The method of claim 2, wherein one or more of the observable features comprise aggregate metrics computed from transaction data for a predefined time period that is defined in relation to a time of the current transaction.

4. The method of claim 3, wherein obtaining a training set of data samples comprises, for a given data sample:

obtaining transaction data for a current transaction, the transaction data comprising an identifier for a uniquely identifiable entity; and

obtaining look-up data for the uniquely identifiable entity;

wherein the obtained transaction data is used to derive the first feature set and the obtained look-up data is used to derive the second feature set.

5. The method of claim 4, wherein obtaining a training set of data samples further comprises, for the given data sample:

obtaining transaction data for the uniquely identifiable entity for the temporal window; and

computing one or more aggregated metrics from the transaction data,

wherein the one or more aggregated metrics are used to derive the first feature set.

6. The method of claim 4 or claim 5, wherein obtaining a training set of data samples further comprises, for the given data sample:

obtaining the historical transaction data for the uniquely identifiable entity, the historical transaction data comprising transaction data that is outside the temporal window; and

computing one or more aggregated metrics from the historical transaction data,

wherein the one or more aggregated metrics are used to derive the second feature set.

7. The method of any one of the preceding claims, wherein the second feature set comprises metadata associated with a corresponding uniquely identifiable entity.

8. The method of any one of the preceding claims, wherein the supervised machine learning system comprises a binary classifier that outputs a value within a predefined range representing a likelihood of an anomaly and wherein the labels indicating the absence or presence of an anomaly comprise two numeric values representing a binary output.

9. The method of any one of the preceding claims, wherein the supervised machine learning system comprises an ensemble system based on a set of decision trees.

10. The method of any one of claims 1 to 8, wherein the supervised machine learning system comprises a recurrent neural network.

11. A method of detecting anomalies within transaction data, comprising:
 - receiving transaction data for anomaly detection;
 - generating a first set of features based on the received transaction data;
 - determining a uniquely identifiable entity associated with the received transaction data;
 - obtaining ancillary data associated with the uniquely identifiable entity;
 - generating a second set of features based on the obtained ancillary data;
 - generating input data for a supervised machine learning system based on the generated first and second set of features, the supervised machine learning system being trained based on a training set comprising historical data samples and synthetic data samples, the synthetic data samples being generated by combining features from the first and second sets of features from the historical data samples that respectively relate to two different ones of the set of uniquely identifiable entities, the synthetic data samples being assigned a label indicating a presence of an anomaly for training and unlabelled data samples within the historical data samples being assigned a label indicating an absence of an anomaly for training;
 - applying the supervised machine learning system to the input data, the supervised machine learning system comprising a binary classifier configured to output a value indicative of a presence of an anomaly; and
 - selectively labelling the received transaction data based on the value output by the supervised machine learning system.
12. The method of claim 11, wherein selectively labelling the received transaction data comprises selectively flagging the uniquely identifiable entity associated with the received transaction data based on the value output by the supervised machine learning system.
13. The method of claim 11 or claim 12, comprising:
 - generating control data to control whether at least one transaction within the transaction data is accepted or denied based on the value output by the supervised machine learning system.
14. The method of any one of claims 11 to 13, wherein obtaining ancillary data associated with the uniquely identifiable entity comprises:
 - obtaining historical transaction data for the uniquely identifiable entity, the historical transaction data comprising transaction data that is outside a temporal window defined in relation to the received transaction data; and
 - computing one or more aggregated metrics from the historical transaction data,

wherein the one or more aggregated metrics are used to generate the second feature set.

15. The method of any one of claims 11 to 14, wherein the transaction data is received from a point-of-sale device in relation to a transaction to be approved, and wherein selectively labelling the received transaction data comprises approving the transaction response to the value output by the supervised machine learning system being below a predefined threshold.