CLAIMS

1. A method for detecting anomalies in double-party interaction data, comprising:
representing interactions between parties of a first group and parties of a second group as a graph, wherein each interaction between a first party of the first group and a second party of the second group is represented by an edge between a respective first node representing the first party and a respective second node representing the second party and wherein information about the first party is assigned to the first node as node attribute information, information about the second party is assigned to the second node as the node attribute information and information about the interaction is assigned to the edge as edge attribute information;
processing the graph by a graph convolutional neural network having an autoencoder structure;
deriving anomaly scores for interactions, parties of the first group and parties of the second group from a reconstruction loss between the graph and an output of the graph convolutional neural network in response to the graph including at least a loss between the edge attribute information and edge attribute information reconstructed by a decoder of the graph convolutional neural network; and
detecting anomalies based on the anomaly scores,
wherein the method further comprises:
training the graph convolutional neural network using a plurality of training data elements, wherein each training data element of the plurality of training data elements comprises a training graph; and
determining the anomaly scores by sampling sub-graphs of the training graph.

2. The method of claim 1, wherein the reconstruction loss includes a loss between the node attribute information and node attribute information reconstructed by the decoder.

3. The method of claim 1 or 2, wherein the reconstruction loss includes a loss between node adjacency information of the graph and node adjacency information reconstructed by the decoder.

4. The method of any one of claims 1 to 3, comprising comparing the anomaly scores with one or more threshold values and detecting an anomaly of an interaction, party of the first group or party of the second group if its anomaly score is above a respective threshold value.

5. The method of any one of claims 1 to 4, wherein the training the graph convolutional neural network is performed by, determining, for each training data element of the plurality of training data elements,
wherein each training data element comprises the training graph, a reconstruction loss between the training graph and the output of the graph convolutional neural network in response to the training graph and the output of the graph convolutional neural network in response to the graph including at least the loss between the edge attribute information and the edge attribute information reconstructed by the decoder of the graph convolutional neural network, and
adapting the neural network to reduce an overall loss including the reconstruction losses determined for the training data elements.

6. The method of claim 5, comprising, for each training data element, sampling the sub-graphs of the training graph, for each sampled sub-graph setting a reconstruction target to the sampled sub-graph and expanding the sub-graph by including nodes and edges connected to the sampled sub-graph and computing a reconstruction between the reconstruction target and the output of the graph convolutional neural network in response to the expanded sub-graph, wherein the overall loss includes the reconstruction losses determined for the sub-graphs.

7. The method of any one of claims 1 to 6, wherein the first group of parties are customers and the second group of parties are service providers.

8. The method of any one of claims 1 to 7, wherein the interactions are transactions between the first group of parties and the second group of parties.

9. The method of any one of claims 1 to 8, comprising detecting fraud based on the detected anomalies.

10. The method of claims 9, comprising checking, for each detected anomaly, whether there has been fraud.

11. The method of any one of claims 1 to 10, further comprising utilizing the anomaly score as an input to a human-in-the-loop actioning system and/or an automatic actioning system.

12. A server computer comprising a radio interface, a memory interface and a processing unit configured to perform the method of any one of claims 1 to 11.

13. A computer program element comprising program instructions, which, when executed by one or more processors, cause the one or more processors to perform the method of any one of claims 1 to 11.

14. A computer-readable medium comprising program instructions, which, when executed by one or more processors, cause the one or more processors to perform the method of any one of claims 1 to 11.