

5

10

15

Monitoring system and method for the electronic monitoring of financial transactions

Description

20 The proposed solution in particular relates to a monitoring system for the electronic monitoring of financial transactions.

The electronic monitoring of financial transactions increasingly uses means of machine learning in order to support for example financial institutions in the automated detection of anomalies in financial transactions. Anomaly here is understood to be a deviation from a predetermined set of rules which is indicative of a fraudulent financial transaction, in particular of at least one financial transaction utilized for money laundering. Electronic monitoring systems are designed to automatically filter out that part of several thousand financial transactions per day which is classified as suspicious. For a financial transaction classified as suspicious an alarm event will then be generated in order to be able to further check the respective financial transaction and possibly report it to a regulatory and/or investigative authority.

Although in this connection elements of artificial intelligence (AI) are increasingly gaining importance, solutions for an efficient, inexpensive and accurate detection of anomalies in financial transactions are hardly known. This is where the proposed solution comes in.

In particular, there is proposed a monitoring system for monitoring financial transactions which electronically checks financial transactions between at least one payer and at least one payee (or corresponding data on such financial transactions) for the occurrence of at least one anomaly. Anomaly here is understood to be a deviation from a particular set of rules, which is indicative of a fraudulent financial transaction, in particular of a financial

transaction utilized for money laundering and hence must be classified as suspicious due to regulatory requirements. On detection of at least one such anomaly, the proposed monitoring system is configured to generate an alarm event and hence an associated electronic alarm signal, which for example triggers an alarm message on a (stationary or
 5 mobile) computer device of the monitoring system and/or on a computer device of a user system coupled with the monitoring system. The proposed monitoring system comprises a tokenizer and a language model for monitoring the financial transactions. The tokenizer here is provided to generate a token sequence from an input sequence representative of a financial transaction. The language model in turn is provided to check the occurrence of an
 10 anomaly in one financial transaction or in several financial transactions for generating an alarm event on the basis of one token sequence or several token sequences.

Consequently, the proposed solution utilizes a language model which, however, here is not provided for processing natural language, but for processing tokenized input sequences for
 15 financial transactions. The proposed solution thus makes use of a language model which instead of a vocabulary for natural language is based on a vocabulary for data on financial transactions, in order to on this basis detect in how far anomalies are present in monitored financial transactions. These anomalies here for example appear as deviations in an input sequence evaluated for a financial transaction or a dataset corresponding herewith as
 20 compared to a setup for the respective financial transaction expected of the language model.

An import sequence here contains transaction characteristics to be distinguished from each other for each financial transaction. In a string handed over to the tokenizer as an input
 25 sequence, these transaction characteristics are treated in the manner of words of a sentence in order to therefrom generate input tokens for the language model. Tokenizing an input sequence thus is effected with reference to the transaction characteristics, which were learned in advance in a machine-based way or manually. A transaction characteristic for example can be an account number, a bank code (known as a "Business Identifier
 30 Code", or BIC for short), a date, a time, or an indicator for a type of transaction (for example transfer or withdrawal).

In one embodiment, the tokenizer is adapted and provided to mask, in a token sequence, at least one input of the input tokens generated from the input sequence, before the token
 35 sequence is handed over to the language model. Here, the monitoring system then is furthermore configured to predict an output token for the at least one masked input token by means of the language model. In this connection, it can then for example furthermore

be provided that for the detection “direction of an anomaly” the monitoring system checks in how far (in particular whether) the at least one predicted output token corresponds with the associated input token. Thus, for the partly masked token sequence of a financial transaction a prediction of the masked part is made via the language model. When the language model cannot predict the masked input token(s), so that the output token does not correspond with the original input token, which has been masked previously, it is inferred on the basis of the language model trained with reference to admissible financial transactions that a financial transaction or a money movement connected therewith is suspicious and hence an anomaly exists, which - possibly by taking account of further criteria - leads to the triggering of an alarm event.

In this connection it can also be provided that a predetermined amount, in particular percentage share of input tokens always is masked, and the monitoring system is configured to check for the detection of an anomaly in how far predicted output tokens of an output sequence for the financial transaction correspond to associated input tokens (i.e. corresponding with respect to their position in the input sequence). When the masked input tokens of one financial transaction or of several financial transactions cannot correctly be predicted by the language model, this speaks for one suspicious financial transaction or for several suspicious financial transactions.

20

Alternatively or additionally it can be provided that for the detection of an anomaly in a financial transaction the monitoring system is configured to compare the financial transactions with reference to at least one reference input token of the token sequence associated to the financial transaction with at least one cluster, which is formed on the basis of token sequences for financial transactions of reference input tokens corresponding with each other. Consequently, in this embodiment data on a financial transaction are compared with one or more clusters by means of the language model and output tokens generated thereby, and it is checked in how far the respective financial transaction can be classified as being associated to a cluster. Here, use is made of the fact that it has been detected that with a specific, corresponding transaction characteristic and hence a reference input token resulting therefrom, financial transactions can be clustered by a language model. However, when a financial transaction with an identical reference input token lies outside a cluster formed in this way beyond a permitted degree, this speaks for a suspicious financial transaction. It can be observed for example that comparable financial transactions always are carried out with accounts of a particular merchant category. For example, a bakery's account typically involves a plurality of payment flows, each involving comparatively small amounts from a plurality of different private accounts during the

bakery's business hours. A financial transaction which represents a payment flow with a much larger amount as compared to an average value outside the bakery's business hours, would thus lie outside a cluster formed for a bakery and would thus be classified as suspicious - possibly by taking account of further criteria.

5

With reference to the exemplary embodiment explained above, the reference input token, which is used for clustering, can each be an input token of a token sequence, which is representative of a merchant category code (MCC) contained in the input sequence.

- 10 For example, it is provided that the monitoring system is configured to determine a data point for the comparison with the at least one cluster with the language model with reference to the at least one reference input token for the financial transaction, and to determine the distance of said data point to a center of the at least one cluster. In such a procedure for example multidimensional, in particular high-dimensional output vectors from
- 15 the language model are subjected to a dimension reduction and subsequently are supplied to a cluster function (i.e., a clustering algorithm) in order to generate, for an input sequence representing a financial transaction, a data point in a two-dimensional or three-dimensional space whose distance to a center of different trained clusters can be determined. For the detection of an anomaly, this determined distance is compared with a stored threshold
- 20 value. When the distance to a particular cluster or to several clusters ultimately is too large, a suspicious financial transaction is inferred and an alarm event is generated.

- In a possible development it can also be provided that different threshold values are stored for different clusters. For clusters of a particular first reference input token a larger scatter
- 25 of the data points for financial transactions can be admissible than in a cluster for financial transactions with a second reference input token different therefrom.

- For generating an alarm event, both of the above-mentioned criteria can also be taken into account in principle. It can be provided, for example, that for the detection of an anomaly it
- 30 is on the one hand checked in how far (whether) at least one predicted output token corresponds with an associated masked input token, and on the other hand it is checked in how far a financial transaction can be assigned to a particular cluster of financial transactions on the basis of an output token sequence. For example, when the language model cannot correctly predict the input tokens masked for a financial transaction, and
- 35 when the financial transaction cannot safely be assigned to a particular cluster by using the language model, this speaks for an anomaly and hence for a suspicious financial transaction which leads to an alarm event. When there possibly is evidence of an anomaly

in only one of the two criteria mentioned above, the same can be weighted differently in order to thereby decide as to whether or not the presence of one of the two criteria already leads to the generation of an alarm event.

- 5 The monitoring system based on a language model for example can be a transformer machine learning model (hence a machine-based transformer learning model). In particular, this can be a transformer machine learning model trained on the basis of input sequences for financial transactions and in particular a bidirectionally trained transformer machine learning model. For example, a transformer machine learning model
- 10 bidirectionally trained with Bidirectional Encoder Representations from Transformers, short: BERT, or another bidirectional Encoder Transformer was found to be particularly advantageous, in order to thereby reliably and with a comparatively manageable calculation effort infer anomalies in monitored financial transactions.
- 15 In principle, the alarm event can be provided for output on a computer device of the monitoring system and/or on a computer device of a user system connected to the monitoring system. Thus, the monitoring system can be implemented for example on the server side, in particular in the cloud. The output of the generated alarm event in turn is effected on a local user system, for example of a financial institution. Alternatively, an on-
- 20 premises monitoring system can be provided. In particular a computer device of a user system for example can comprise a mobile terminal, such as a laptop or a mobile phone.

The proposed solution furthermore relates to a method for the electronic monitoring of financial transactions between at least one payer and at least one payee for the occurrence

25 of at least one anomaly. It here is proposed to generate a token sequence from an input sequence representative of a financial transaction via one tokenizer each and to check the occurrence of an anomaly in one or more financial transactions by using a language model on the basis of one token sequence or several token sequences, in order to generate an alarm event in response to at least one anomaly detected during the check.

30

An embodiment of a proposed method in particular can be realized with an embodiment of a proposed monitoring system. Correspondingly, the advantages and features of embodiments of a proposed monitoring system as explained above and below also apply for embodiments of a proposed method, and vice versa.

35

In particular, in one embodiment of the proposed method it can consequently be provided that the tokenizer in a token sequence masks at least one input token of the input tokens

generated from the input sequence, before the token sequence is handed over to the language model, in order to subsequently predict the at least one masked input token by using the language model and to check an output token generated for this purpose from the language model, in how far the same corresponds with the unmasked input token.

5

Likewise, for the detection of an anomaly it can alternatively or additionally be provided that it is checked with reference to at least one reference input token for a financial transaction in how far this financial transaction can be assigned to a cluster of comparable financial transactions by utilizing the language model.

10

Furthermore, the proposed solution relates to a computer program product comprising commands which on execution by at least one processor cause the at least one processor to execute an embodiment of a proposed method.

15 In addition, the proposed solution also relates to a method for training a language model for monitoring financial transactions. Consequently, the proposed solution not only relates to the use of the language model trained already for the detection of anomalies in financial transactions. Rather, the focus also is on how a corresponding language model can be trained. In this connection it is proposed that a corresponding training method comprises
20 at least the following steps:

- providing data on individual financial transactions, typically each in the form of a string,
- generating at least one set of input sequences for the financial transactions, wherein each input sequence contains transaction characteristics to be distinguished from
25 each other (typically predefined and provided to the language model as an input);
- tokenizing the input sequences with reference to the transaction characteristics by using a tokenizer for generating input tokens for each input sequence;
- training the language model on the basis of the input tokens and the input sequences by masking at least part of the input tokens, in order to provide a language model
30 trained for data on financial transactions, by means of which the occurrence of an anomaly in one or more financial transactions can be detected on the basis of one token sequence or several token sequences.

In the proposed method the training of the language model thus is provided with the
35 objective that the language model correctly determines the masked input tokens from the context provided and hence checks in how far output tokens generated by means of the language model correspond with input tokens which previously were provided to the

language model only in masked form. According to this aspect of the proposed solution there is hence utilized a (possibly pre-trained) language model in order to learn a “vocabulary” on the basis of tokenized input sequences for financial transactions.

- 5 In principle, it can be provided in this connection that the input tokens each are assigned to vectors of real numbers by means of word embedding. Before that, the tokenizer can convert the transaction characteristics into numerical input tokens. From the perspective of the language model, for example, when a dataset and hence a string which characterizes a financial transaction can be understood as a word, then a set of input sequences for a particular account and/or a particular merchant category describes a text containing words characteristic therefor. As already explained above, this means that, for example, the financial transactions underlying the at least one set of input sequences can be assigned to an account of a particular merchant category code during the training of the language model. A set of input sequences then for example forms a “text example” for a particular type of account, which can be identified with reference to a reference input token (typically preceding an input sequence). Between individual tokenized input sequences for the respective account, separation or separating tokens can each be provided, for example.

For example, two input sequences for a particular account can be constructed as follows:

20

(1)

Trx1 Hp1 Cp019 2023-10-27:00:00:00 300EUR CASH DEBIT
ATM_WITHDRAWAL DE MCC5892

25

(2)

Trx2 Hp1 Cp670 2023-10-30:01:10:08 600EUR ACH CREDIT ACH_TRANSFER
US MCC5892

- 30 With reference to predefined transaction characteristics, these input sequences can be represented as follows:

(3)

0D 0H 0M 0S AMOUNT2 CASH DEBIT ATM_WITHDRAWAL DE

35

(4)

3D 1H 10M 8S AMOUNT3 ACH CREDIT ACH_TRANSFER US

For example, the time ("2023-10-27:00:00:00") of the earliest financial transaction (1) is set as zero point and for the date and time specifications for the second financial transaction merely the difference ("3D 1H 10M 8S") is used.

5

Subsequently, the input sequences for the financial transactions (1) and (2) are transferred into a numerical representation on the basis of the formats (3) and (4), as follows:

(5)

10 0 31 54 67 120 129 141 143 286

(6)

3 32 65 74 121 132 140 148 232

15 Assuming that the corresponding account refers to a merchant category code MCC with the value "523", an input token sequence can be represented therewith by the tokenizer for this account as follows:

20 [CLS] 523 [SEP] 3 32 65 74 121 132 140 148 232 [SEP] 0 31 54 67 120 129 141
143 286 [SEP] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD]
[PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD]
[PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD]
[PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD]
[PAD]

25

The token [CLS] here forms a usual classification token. [SEP] refers to a separation or separating token which separates two input sequences from each other. [PAD] designates a padding token in order to bring a set of input token sequences for an account to a size uniform for all training data.

30

What can also be part of the tokenized transaction characteristics and hence of the numerical input tokens is a categorization parameter for the transaction volume of the respective financial transaction. Thus, an IBAN of a counterparty for each account under consideration, for which a financial transaction is observed, generally is different. If an IBAN
35 were handed over as a transaction characteristic, this would make it difficult to transfer learned characteristics from one account to another account, since each account and its transaction characteristics observed would be associated with a different set of IBANs and

hence input tokens. Against this background a categorization or sorting according to transaction volume is provided in one embodiment. Here, a set of input sequences will be provided to the language model to be trained for each account, which contains a value for a categorization parameter. The values for this categorization parameter depend on the transaction volume of the respective financial transaction with the account of the counterparty. A first value for this categorization parameter for example is always set to be identical for the account with the largest transaction volume. For example, it can be provided in this connection that for the account with the largest transaction volume there is always provided an input token "0" for the categorization parameter, and in dependence on the size of the occurring transaction volumes further input tokens with ascending values are set, e.g. an input token "1" for the second largest transaction volume, etc. This permits a better generalization from one account to the next account for the language model.

For training the language model, a plurality of sets of input sequences can be generated, which are based on accounts with identical and different merchant category codes. On this basis, output vectors generated as output by using the language model - after an optional dimension reduction and by means of a cluster function - clusters can be generated, which are each representative of financial transactions of a merchant category code. These clusters can be utilized for a future comparison in how far a financial transaction with a particular merchant category code to be checked with the trained language model can be assigned to a corresponding cluster for this merchant category code or possibly deviates therefrom to such an extent that an anomaly must be assumed.

In principle, the language model to be trained can be a transformer machine learning model, in particular a transformer machine learning model bidirectionally pre-trained with BERT, a BERT-based encoder transformer model or another bidirectional encoder transformer model.

The attached Figures by way of example illustrate possible embodiments of the proposed solution.

In the drawings:

Figure 1 shows a schematic representation of the monitoring system and its monitoring of financial transactions;

- Figure 2 shows a set of input sequences for various financial transactions by illustrating different transaction characteristics for an individual transaction dataset;
- 5 Figure 3A shows the procedure when training a transformer machine learning model bidirectionally pre-trained on the basis of input sequences as shown in Figure 2;
- 10 Figure 3B shows the procedure for checking financial transactions for possible anomalies on the basis of the trained language model;
- Figure 4A shows a visualization of different clusters for financial transactions determined by means of the trained language model;
- 15 Figure 4B shows a visualization for a data point determined by means of the language model for a suspicious financial transaction by illustrating distances of this data point to different clusters;
- 20 Figure 5A shows a schematic representation for the assessment of the relevance of financial transactions on the basis of transaction characteristics or input tokens defined therewith;
- 25 Figure 5B shows a schematic representation for the processing of the relevance of different transaction characteristics of a financial transaction on generation of an alarm event.

Figure 1 schematically shows an electronic monitoring system 1 which can be used for electronically monitoring financial transactions by means of a language model. Data on a plurality of financial transactions here are provided to the monitoring system 1, for example

30 from one financial institution or several financial institutions, or the monitoring system 1 is used at a financial institution. Figure 1 schematically shows financial transactions T1 to T4 for the account of a payee E, hence a beneficiary. With each financial transaction T1 to T4 the payee E receives payment flows of varying amounts from different payers S1, S2, S3. By means of the monitoring system 1, the financial flows are checked on the basis of the

35 financial transactions T1 to T4, in order to electronically and automatically identify suspected cases in which an anomaly in the respective financial transaction and hence for example a violation of regulations concerning money laundering is assumed.

When an anomaly is detected in one of the financial transactions, the monitoring system 1 automatically generates an alarm event. This alarm event then for example leads to the output of an alarm message in order to indicate a suspected case with at least one of the financial transactions T1 to T4. The output of such an alarm message then for example is
 5 effected on a computer device of a user system 2 coupled with the monitoring system 1 in connection with a fight against money laundering (anti-money laundry).

In one embodiment of the proposed solution it is provided to provide data on each financial transaction T1 to T4 for a trained language model of the monitoring system 1 as an input
 10 sequence SQ with transaction characteristics TC1 to TC5 to be distinguished from each other. An input sequence SQ corresponding to Figure 2 for example includes a transaction text TT and hence a string which is composed of the plurality of transaction characteristics TC1 to TC5.

Via the sum of the transaction characteristics TC1 to TC5 a financial transaction T1 to T4 is determined unambiguously. For example, a first transaction characteristic TC1 includes a statement on the account of the payee E. A value for the transaction characteristic TC2 in turn represents the kind of financial transaction and for example indicates whether it is an ATM withdrawal or a bank transfer. A further transaction characteristic TC3 for example
 20 indicates the account number of the payer S1, S2 or S3. The transaction characteristic TC4 for example identifies the bank code of the account of the payer S1, S2, S3. A transaction characteristic TC5 in turn contains an information regarding the amount of money. Further transaction characteristics in turn indicate the date and time of the financial transactions.

For a training of a language model of the monitoring system 1 a plurality of input sequences SQ can be provided for an account in the manner of words which should form the basis for a vocabulary to be learned. From financial transactions for an individual account a set of several financial transactions therefore is arbitrarily selected in order to provide a transaction dataset TS and hence a set of input sequences SQ for the training.
 30

Both for the training of the language model and for its future use an input sequence SQ each is tokenized. The individual transaction characteristics TC1 to TC5 here are converted into numerical input tokens w_1 to w_5 .

For example, two input sequences for the one account can be constructed as follows:

(1)

Trx1 Hp1 Cp019 2023-10-27:00:00:00 300EUR CASH DEBIT
ATM_WITHDRAWAL DE MCC5892

5

(2)

Trx2 Hp1 Cp670 2023-10-30:01:10:08 600EUR ACH CREDIT ACH_TRANSFER
US MCC5892

With reference to the predefined transaction characteristics TC1 to TCn, these input

10 sequences can be represented as follows:

(3)

0D 0H 0M 0S AMOUNT2 CASH DEBIT ATM_WITHDRAWAL DE

15

(4)

3D 1H 10M 8S AMOUNT3 ACH CREDIT ACH_TRANSFER US

For example, the time ("2023-10-27:00:00:00") of the earliest financial transaction (1) is set
as zero point and for the date and time specifications for the second financial transaction

20 merely the difference ("3D 1H 10M 8S") is used.

Subsequently, the input sequences for the financial transactions (1) and (2) are transferred
into a numerical representation for input tokens w_1 to w_n on the basis of the formats (3) and
(4), as follows:

25

(5)

0 31 54 67 120 129 141 143 286

(6)

30

3 32 65 74 121 132 140 148 232

Assuming that the corresponding account relates to a merchant category code MCC with
the value "523", an input token sequence can be represented therewith by the tokenizer for
this account as follows:

35

5 [CLS] 523 [SEP] 3 32 65 74 121 132 140 148 232 [SEP] 0 31 54 67 120 129 141
 143 286 [SEP] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD]
 [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD]
 [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD]
 [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD] [PAD]
 [PAD]

10 The token [CLF] here forms a usual classification token. [SEP] refers to a separation or
 separating token which separates two input sequences from each other. [PAD] designates
 a padding token in order to bring a set of input token sequences SQ for an account to a
 size for a transaction data set TS which is uniform for all training data.

15 What can also be part of the tokenized transaction characteristics TC1 to TCn and hence
 the numerical input tokens is a categorization parameter for the transaction volume of the
 respective financial transaction T1 to T4. Thus, an IBAN of a payee E for each account
 under consideration, for which a financial transaction T1 to T4 is observed, generally is
 different. If an IBAN were handed over as a transaction characteristic, this would make it
 difficult to transfer learned characteristics from one account to another account, as with
 20 each account and its transaction characteristics TC1 to TCn observed a different set of
 IBANs and hence input tokens w_1 to w_n would be supplied. Against this background a
 categorization or sorting according to transaction volume is provided in one embodiment.
 Here, a value for a categorization parameter is associated to each account as part of a set
 of input sequences SQ, which in the example of Figure 1 are provided to the language
 25 model to be trained for a payer S1 to S3, hence to an account of a counterparty. The values
 for this categorization parameter depend on the transaction volume of the respective
 financial transaction T1 to T4 with the respective account (of the counterparty). A first value
 for this categorization parameter for example is always set to be identical for the account
 of a payer S1 to S3 with the largest transaction volume. For example, it can be provided in
 30 this connection that for the account with the largest transaction volume there is always
 provided an input token with the value "0" for the categorization parameter, and in
 dependence on the size of the occurring transaction volumes further input tokens with
 ascending values are set, e.g. an input token "1" for the second largest transaction volume,
 etc. This permits a better generalization from one account to the next account for the
 35 language model.

The training of the language model used in the monitoring system 1 for monitoring financial transactions T1 to T4 is illustrated in more detail with reference to Figure 3A. Here, there is provided a language model 31 with bidirectional transformer encoder. Input sequences SQ or a set of input sequences SQ here is/are supplied to a tokenizer 10. On the basis of the transaction characteristics TC1 to TC5 of each input sequence SQ for the financial transactions, this tokenizer 10 generates a token sequence with the input tokens w_1 to w_5 .

For training the language model 31, a particular amount of input tokens w_1 to w_5 now is masked, in the example of Figure 3 illustrated by way of example by replacement of the input token w_5 by a masked token [MASK]. The token sequence then is supplied to a word embedding 30, in order to prepare the token sequence for the language model 31. This preparation here includes the addition of the classification token [CLS], of the separating tokens [SEP] and of the padding tokens [PAD], as explained above.

A token sequence then passes through a transformer encoder plane of the language model 31. For a partly masked token sequence, the language model 31 in a manner known per se for BERT generates multidimensional output vectors O1 to O5, from which after a classification 32 and subsequent vocabularization 33 output tokens w'_1 to w'_5 are obtained. The output vectors O1 to O5 correspondingly correspond with the respective input tokens w_1 to w_5 . The output vectors O1 to O5 consequently contain the context information of the respective input token w_1 to w_5 , so that it can be learned in what context a particular transaction characteristic occurs. In the classification 32 one or more classification layers are passed through, which generate probabilities for various classes in a classification task, so that subsequently output tokens for masked input tokens, such as an output token w'_4 for the masked input token w_4 , are inferred on the basis of a probability distribution and of the context learned already. Consequently, the language model 31 is trained with a plurality of sets of input sequences SQ with the objective to infer masked input tokens such as the input token w_4 on the basis of the context, hence of unmasked input tokens w_1, w_2, w_3, w_5 . For example, 15% of the input tokens w_1 to w_5 can be randomly selected for masking. Of these, i) 80% are replaced by a masked token [MASK] corresponding to Figure 3A, ii) 10% are replaced with a random input token, and iii) 10% are left unchanged. By comparing the generated output tokens w'_1 to w'_5 with the token sequence of the input tokens w_1 to w_5 , the language model 31 is trained to correctly predict input tokens and hence to infer transaction characteristics TC1 to TC5.

When the language model is trained with corresponding training data, a trained language model 34 corresponding to Figure 3B is available for the efficient monitoring of financial

transactions. Here, for each financial transaction T1 to T4 an input sequence SQ is provided to the tokenizer 10 individually or as part of a set of several input sequences SQ for an account. A particular amount of the input tokens w_1 to w_5 here is masked as well. By passing through the transformer encoder plane of the trained language model 34 high-dimensional output vectors O1 to O5 are generated, by means of which output vectors are predicted for masked input tokens w_4 . When predicted output tokens and hence an output token sequence deviate from the actual token sequence w_1 to w_5 , this suggests an anomaly in the respective financial transaction. This means, it is checked in other words in how far previously masked input tokens are correctly predicted for transaction characteristics on the basis of the trained language model 34. When this is not the case, the transaction characteristics TC1 to TC5 of a financial transaction strikingly deviate from the normal case. Hence the respective financial transaction appears to be suspicious and triggers an alarm event, in order to check the same further in detail by a user or a monitoring body, for example at a financial institution. A corresponding anomaly detection 35 is provided in Figure 3B.

The anomaly detection 35 possibly can include in addition that after a dimensional reduction of the high-dimensional output vectors O1 to O5 a data point is determined for a financial transaction and is compared with clusters for the previously considered financial transactions, which were generated when training the language model 31.

Corresponding to the visualization of Figure 4 it can thus be shown that financial transactions, which are classified as a reference input token for an account with a particular merchant category code MCC, can be assigned to different clusters C1 to C9 specific for merchant category codes by using the language model 31. However, when a data point A for a financial transaction corresponding to Figure 4B does not lie in the vicinity of one of these clusters C1 to C9, this also speaks for an anomaly in the financial transaction which belongs to the data point A. For example, in the visualized example of Figure 4B, the data point A with a distance d_1 to a center M1 of the one cluster C8 and with a distance d_2 to a center M2 of another cluster C5 each is spaced apart so much that the distance d_1 lies above a threshold value specified for the cluster C8 and the distance d_2 lies above a threshold value specified for the cluster C5.

Thus, when it is noted for example for a financial transaction by means of the trained language model 34 in connection with the anomaly detection 35 of Figure 3B that the language model 34 does not completely correctly predict the token sequence w_1 to w_5 and, moreover, a data point A determined therefor cannot be assigned to a previously

determined cluster C5 or C8 for similar financial transactions, this speaks for an anomaly and hence a suspicious financial transaction, which triggers an alarm event.

To make a correspondingly triggered alarm event more easily assessable in connection with the assessment of a further check by a human user, the embodiment of Figure 3B also provides a relevance assessment 36 after the anomaly detection 35. The criteria which were relevant as suspicious for the language-model-based assessment of a financial transaction here are prepared for a human user on output of the alarm event and are comprehensibly visualized.

For example, each financial transaction T1 to T4 can be represented as a sum of transaction characteristics or corresponding transaction tokens 500, 501, which are assigned to different groups 50 and 51. For example, the transaction tokens 500 jointly represent a payment type 50 and the transaction tokens 501 represent a country code 51. Since for example in the financial transaction T1 one of the transaction tokens 500 for the payment type 50 - with regard to the detection of an anomaly - is classified as highly relevant, the financial transaction is provided with an importance score of 1. On the other hand, the financial transaction T3 receives an importance score of 2, while the two further illustrated items of information T2 and T4 are classified with an importance score of 3. In a visualization of the different financial transactions T1 to T4, the financial transaction T1 thus is emphasized on a user system 2 as possibly more relevant for the more exact consideration in response to the occurrence of an anomaly. Thus, the trained language model 34 provides for a token-based explainability for the classification of various financial transactions as potentially suspicious or unsuspicious.

This principle is also based on a visualization corresponding to Figure 5B. While the visualization in connection with the relevance assessment 36 corresponding to Figure 5A uses a token-based approach to visualize differences across several financial transactions T1 to T4, a visualization corresponding to the procedure shown in Figure 5B is intended to visualize, for a user of the user system 2, the different transaction characteristics within a financial transaction T1, which have led to an assessment of this financial transaction T1 as suspicious and hence as the detection of an anomaly. For example, when the output tokens of the trained language model 34 - predicted on the basis of transaction tokens 500 and 501 and hence the corresponding transaction characteristics - suggest that there are significant deviations from financial transactions to be classified as permissible, especially with regard to payment type and country code, these deviations can be weighted and visualized accordingly.

For example, after several similarly suspicious financial transactions have occurred, a user may be notified that the alarm event triggered for an account in particular on the basis of the financial transaction T1 was in part due to the fact that repeatedly in quick succession smaller amounts were transferred to a particular account of a payee E, a large part of the transaction volume each was part of a round-trip transaction and at least partly was attributable to cross-border financial transactions.

By means of a likewise token-based visualization, the explainability for a user of the monitoring system 1 can be improved considerably. This offers an additional advantage in the technique improved in connection with the proposed solution for the efficient electronic, automated monitoring of financial transactions on the basis of a language model which can be implemented with comparatively low processing power and high speed.

List of reference numerals

1	monitoring system
10	tokenizer
2	user system
30	word embedding
31	pre-trained language model / transformation encoder
32	classification
33	vocabularization
34	trained language model
35	anomaly detection and clustering
36	relevance assessment
50	payment type
500	transaction token for payment type
501	transaction token for country code
51	country code
A	data point for anomaly
C1 - C9	cluster
d1, d2	distance
E	payee
I	importance score
M1, M2	cluster center
O1, O2, O3, O4, O5	output vector
S1, S2, S3	payer
SQ	input sequence
T1, T2, T3, T4	financial transaction
TC1 - TC5	transaction characteristic
TS	transaction dataset (set of input sequences)
TT	transaction text
$w'_1, w'_2, w'_3, w'_4, w'_5$	output token
w_1, w_2, w_3, w_4, w_5	input token

Claims

1. A monitoring system for monitoring financial transactions (T1, T2, T3, T4), wherein the monitoring system (1) electronically checks financial transactions (T1, T2, T3, T4) between at least one payer (S1, S2, S3) and at least one payee (E) for the occurrence of at least one anomaly, and the monitoring system (1) is configured to generate an alarm event on detection of at least one anomaly,

characterized in that

the monitoring system (1) comprises a tokenizer (10) and a language model (34) for monitoring the financial transactions (T1, T2, T3, T4), wherein the tokenizer (10) is provided to generate a token sequence (w_1 - w_5) from an input sequence (SQ) representative of a financial transaction (T1, T2, T3, T4), and the language model (34) is provided to check the occurrence of an anomaly in one or more financial transactions (T1, T2, T3, T4) for generating an alarm event on the basis of one token sequence (w_1 - w_5) or several token sequences (w_1 - w_5).
2. The monitoring system according to claim 1, **characterized in** that the tokenizer (10) is adapted and provided to mask in a token sequence (w_1 - w_5) at least one input token (w_4) of the input tokens (w_1, w_2, w_3, w_4, w_5) generated from the input sequence (SQ), before the token sequence (w_1 - w_5) is handed over to the language model (34), and the monitoring system is configured to use the language model (34) to predict an output token (w'_4) for the at least one masked input token (w_4).
3. The monitoring system according to claim 2, **characterized in** that the monitoring system is configured to check for the detection of an anomaly in how far the at least one predicted output token (w'_4) corresponds with the associated input token (w_4).
4. The monitoring system according to claim 3, **characterized in** that the tokenizer (10) is adapted and provided to mask a specified amount of input tokens (w_1, w_2, w_3, w_4, w_5) in a token sequence (w_1 - w_5), and the monitoring system is configured to check for the detection of an anomaly in how far predicted output tokens ($w'_1, w'_2, w'_3, w'_4, w'_5$) correspond with associated input tokens (w_1, w_2, w_3, w_4, w_5).
5. The monitoring system according to any of the preceding claims, **characterized in** that for the detection of an anomaly the monitoring system (1) is configured to compare a financial transaction (T1, T2, T3, T4) with reference to at least one reference input

token (w_1) of the token sequence (w_1-w_5) associated to the financial transaction (T1, T2, T3, T4) with at least one cluster (C1-C9), which is formed on the basis of token sequences for financial transactions (T1, T2, T3) with reference input tokens (w_1) corresponding to each other.

5

6. The monitoring system according to claim 5, **characterized in** that the reference input token (w_1) each is an input token of a token sequence (w_1-w_5) which is representative of a merchant category code contained in the input sequence (SQ).

10 7. The monitoring system according to claim 5 or 6, **characterized in** that the monitoring system (1) is configured to check for the detection of an anomaly in how far the respective financial transaction (T1, T2, T3, T4) can be assigned to the at least one cluster (C1-C9).

15 8. The monitoring system according to any of claims 5 to 7, **characterized in** that the monitoring system (1) is configured to determine a data point (A) for the comparison with the at least one cluster (C1-C9) by using the language model (34) with reference to the at least one reference input token (w_1) for the financial transaction (T1, T2, T3, T4), and to determine its distance (d_1 , d_2) to a center (M1, M2) of the at least one
20 cluster (C1-C9).

9. The monitoring system according to claim 8, **characterized in** that the monitoring system (1) is configured to compare the determined distance (d_1 , d_2) with a stored threshold value for the detection of an anomaly.

25

10. The monitoring system according to claim 9, **characterized in** that for different clusters (C1-C9) different threshold values are provided.

11. The monitoring system according to any of claims 8 to 10, **characterized in** that for the
30 determination of the data point (A) the monitoring system (1) is configured to carry out a dimension reduction for multi-dimensional output vectors (O1-O5) obtained from the language model (34) for input tokens (w_1 , w_2 , w_3 , w_4 , w_5) of a financial transaction (T1, T2, T3, T4).

35 12. The monitoring system according to any of claims 3 or 4 and any of claims 5 to 11, **characterized in** that for the detection of an anomaly the monitoring system (1) is configured to on the one hand check in how far the at least one predicted output token (w'_4) corresponds with the associated input token (w_4), and on the other hand to carry out a comparison with the at least one cluster (C1-C9) with reference to the at least
40 one reference input token (w_1).

13. The monitoring system according to any of the preceding claims, **characterized in** that the language model (34) of the monitoring system (1) is a transformer machine learning model, in particular a transformer machine learning model trained on the basis of input sequences (SQ) for financial transactions (T1, T2, T3, T4), and in particular a
5 transformer machine learning model bidirectionally trained on the basis of input sequences (SQ) for financial transactions (T1, T2, T3, T4).
14. The monitoring system according to any of the preceding claims, **characterized in** that the alarm event is provided for output on a computer device of the monitoring system
10 (1) and/or on a computer device of a user system (2) connected to the monitoring system (1).
15. A method for the electronic monitoring of financial transactions (T1, T2, T3, T4) between at least one payer (S1, S2, S3) and at least one payee (E) for the occurrence
15 of at least one anomaly, wherein in response to at least one anomaly detected during the check an alarm event is generated,
- characterized in that**
- 20 from an input sequence (SQ) representative of a financial transaction (T1, T2, T3) a token sequence (w_1 - w_5) each is generated via a tokenizer (10) and on the basis of one token sequence (w_1 - w_5) or several token sequences (w_1 - w_5) the occurrence of an anomaly in one or more financial transactions (T1, T2, T3) is checked for generating an alarm event by using a language model (34).
25
16. A computer program product comprising commands which on execution by at least one processor cause the at least one processor to carry out the method according to claim 15.
- 30 17. A method for training a language model (31) for monitoring financial transactions, comprising at least the following steps:
- providing data (TT) on individual financial transactions (T1, T2, T3, T4);
 - generating at least one set of input sequences (SQ) for the financial transactions, wherein each input sequence (SQ) contains transaction characteristics (TC1-TC5)
35 to be distinguished from each other;
 - tokenizing the input sequences (SQ) with reference to the transaction characteristics (TC1-TC5) by using a tokenizer (10) for generating input tokens (w_1 - w_5) for each input sequence (SQ);

– training the language model (31) on the basis of the input tokens (w_1, w_2, w_3, w_4, w_5) and the input sequences (SQ) by masking at least part of the input tokens (w_1, w_2, w_3, w_4, w_5) in order to provide a language model (34) trained for data on financial transactions (T1, T2, T3), by means of which the occurrence of an anomaly in one or more financial transactions (T1, T2, T3) can be detected on the basis of one token sequence (w_1-w_5) or several token sequences (w_1-w_5).

18. The method according to claim 17, **characterized in** that the input tokens (w_1, w_2, w_3, w_4, w_5) are assigned to vectors of real numbers by means of word embedding.

10

19. The method according to claim 17 or 18, **characterized in** that the tokenizer (10) converts the transaction characteristics (TC1-TC5) into numerical input tokens (w_1, w_2, w_3, w_4, w_5).

20. The method according to claim 19, **characterized in** that the financial transactions (T1, T2, T3, T4) underlying the at least one set of input sequences (SQ) are assigned to an account of a particular merchant category code.

21. The method according to any of claims 17 to 20, **characterized in** that a plurality of sets of input sequence (SQ) are generated, wherein these sets of input sequences (SQ) are underlying accounts with different merchant category codes, and on the basis of multi-dimensional output vectors (O1-O5) generated as output by the language model (31), after a dimensional reduction, clusters (C1-C9) are generated via a cluster function, which each are representative of financial transactions (T1, T2, T3, T4) of a merchant category code.

22. The method according to any of claims 17 to 21, **characterized in** that at least one of the tokenized transaction characteristics (TC1-TC5) relates to a categorization parameter for a transaction volume of the respective financial transaction (T1, T2, T3, T4).

23. The method according to claim 22, **characterized in** that values for the categorization parameter are specified in ascending or descending order starting with a value for that account of a counterparty (S1-S3) which in view of the provided data (TT) on the financial transactions (T1, T2, T3, T4) with an account has the largest transaction volume.

24. The method according to any of claims 17 to 23, **characterized in** that the language model (31) to be trained is a transformer machine learning model, in particular a bidirectionally pre-trained transformer machine learning model.
- 5 25. Use of a monitoring system according to any of claims 1 to 14 for monitoring financial transactions.