Method and System for Assessing the Reputation of a Merchant

BACKGROUND

(0001) Payment fraud is a serious and growing concern for merchants, especially online retailers. Apart from the financial cost, fraud and/or chargebacks also impact merchant reputation and increase customer attrition. Further, excessive chargebacks may result in a merchant being placed on a watchlist known as MATCH – MasterCard Alert to Control High-risk Merchants. Once a merchant is placed on MATCH, it may remain on the watchlist for a period of 5 years. A placement on MATCH may significantly harm a merchant. Not only does it follow a termination of the merchant account, but it also may serve to inform other acquiring banks that there is something seriously wrong with the business. Consequently, being added to a watchlist often results in the end of payment processing capabilities for the merchant as it becomes extremely difficult for a merchant to find a new acquiring bank, on any terms.

(0002) Merchants may be aware of the consequences, but they lack the capital and resources to fight fraudulent transaction and/or chargebacks. Small merchants with low transaction volume are especially at risk as they do not have access to a significant portion of the transactions across the payment network to develop an effective fraud detection solution. Therefore, there is a need to develop a dynamic and reliable system that can precisely assess the merchant risk level.

SUMMARY

(0003) A system and method are described that may evaluate an adaptive risk score and caution an innocent merchant when it is being be attacked by illicit fraudsters or before it is placed on the MATCH watchlist. The system and method may have the contributions:

(0004) • Near-real time: Detailed statistics on the percentage of fraud and chargeback transactions may not be known to the merchant immediately as it could take one to six months for the chargeback process to be completed. Therefore, quite often the merchants may have no idea of these chargebacks and frauds until they are informed by acquiring banks. Also, given the temporal nature of fraudulent activity, a

previously low-risk merchant could become a high-risk merchant for a short-period of time due to a burst of fraud attacks and then drop back to normal levels of frauds and/or chargebacks. The proposed system and method may be dynamic in that it may evaluate the merchant risk level in near real-time and provide early-warning notification to the merchant.

(0005) • Continuous vs Discrete: The system and method may improve upon the traditional approach of categorizing merchants into two groups: low-risk and high-risk. The proposed system and method may be able to assess a continuous risk score proportional to the likelihood of a merchant being risky, effectively addressing the issue of 'shades of gray' encountered in the traditional blacklisting approach. A continuous risk score may offer greater flexibility to a payment processing network as it allows the network to make dynamic pricing decisions (known as interchange optimization). For instance, the network may decide to charge a higher interchange fee from a merchant whose transactions have an elevated risk.

(0006) • Reliability: The system and method may evaluate the merchant risk level based on collective transaction behavior from merchants across the whole payment network. Each transaction from a user may be scored not only based on the information associated with the current merchant, but also based on the information associated with previous merchants where the user transacted. The merchant risk level may be obtained by aggregating the risk scores over all the transactions that occurred at a particular merchant. Therefore, the merchant risk level obtained may be quite reliable which may be especially beneficial to small merchants with low transaction volume as even a few fraudulent transactions can easily push them into a high-risk merchant category.

(0007) Further, the system and method may help payment processing networks make better decision on cross-border transactions. Merchants in some countries such as China and Russia have no obligations to report fraudulent transactions. Using collective intelligence from past user transactions at merchants that report frauds, the system and method may assess the risk level of merchants that don't report frauds.

(0008)

(0009) Thus, the system and method may assess the merchant risk level on a more continuous scale rather than a binary categorization. It may produce a continuous risk score proportional to the likelihood of a merchant being risky, effectively addressing the issue of shades of gray encountered by the traditional blacklisting approach. The continuous risk score feature provides greater flexibility as it allows the payment network to make dynamic pricing decisions (known as interchange optimization) based on the merchant risk level. Using collective intelligence from transactions across the payment network, the system and method may be able to assess the merchant risk level with high accuracy. The system and method may be particularly beneficial to small merchants with low transaction volume as even a few fraudulent transactions can easily put them in the high-risk merchant category. Further, the system and method may help payment processing networks make better decision on cross-border transactions.

BRIEF DESCRIPTION OF THE DRAWINGS

(00010) FIG. 1 may be an illustration of the computers used in the system;

(00011) FIG. 2 may be an illustration of the computing elements that may be part of the system and may execute the method;

(00012) FIG. 3 may be a flowchart of blocks undertaken by the method;

(00013) FIG. 4 may illustrate an example RNN model architecture;

(00014) FIG. 5 may illustrate the collection of data on an hourly basis which is aggregated periodically;

(00015) FIG. 6 may illustrate an initial data set;

(00016) FIG. 7A may illustrate the rotation of training sets;

(00017) FIG. 7B may illustrate the rotation of training sets;

(00018) FIG. 8 may illustrate a point anomaly where a single spike 805 may be seen in the graph;

(00019) FIG. 9 may illustrate a score shift 905 (from low to high) followed by two spikes; and

(00020) FIG. 10 may illustrate a sudden jump in risk score 1005 and the score stays high for a period of time.

Persons of ordinary skill in the art will appreciate that elements in the figures are illustrated for simplicity and clarity so not all connections and options have been shown to avoid obscuring the inventive aspects. For example, common but well-understood elements that are useful or necessary in a commercially feasible embodiment are not often depicted in order to facilitate a less obstructed view of these various embodiments of the present disclosure. It will be further appreciated that certain actions and/or steps may be described or depicted in a particular order of occurrence while those skilled in the art will understand that such specificity with respect to sequence is not actually required. It will also be understood that the terms and expressions used herein are to be defined with respect to their corresponding respective areas of inquiry and study except where specific meaning have otherwise been set forth herein.

SPECIFICATION

(00021) The present disclosure now will be described more fully hereinafter with reference to the accompanying drawings, which form a part hereof, and which show, by way of illustration, specific exemplary embodiments by which the disclosure may be practiced. These illustrations and exemplary embodiments are presented with the understanding that the present disclosure is an exemplification and is not intended to be limited to any one of the embodiments illustrated. The disclosure may be embodied in many different forms and should not be construed as limited to the embodiments set forth herein; rather, these embodiments are provided so that this disclosure will be thorough and complete, and will fully convey the scope of the disclosure to those skilled in the art. Among other things, the present disclosure may be embodied as methods or devices. Accordingly, the present disclosure may take the form of an entirely hardware embodiment, an entirely software embodiment or an embodiment combining software and hardware aspects. The following detailed description is, therefore, not to be taken in a limiting sense.

(00022) The system 100 may include a computer network 102 that links one or more systems and computer components. In some embodiments, the system 100 includes a user computing device 104, a merchant computer system 106, a payment network system 108 and a transaction analysis system 110.

(00023) The network 102 may be described variously as a communication link, computer network, internet connection, etc. The system 100 may include various software or computer-executable instructions or components stored on tangible memories and specialized hardware components or modules that employ the software and instructions in a practical application to manage network nodes for a transaction, as described herein.

(00024) The various modules may be implemented as computer-readable storage memories containing computer-readable instructions (i.e., software) for execution by one or more processors of the system 100 within a specialized or unique computing device. The modules may perform the various tasks, steps, methods, blocks, etc., as described herein. The system 100 may also include both hardware and software applications, as well as various data communications channels for communicating data between the various specialized and unique hardware and software components.

(00025) Networks are commonly thought to comprise the interconnection and interoperation of hardware, data, and other entities. A computer network, or data network, is a digital telecommunications network which allows nodes to share resources. In computer networks, computing devices exchange data with each other using connections, e.g., data links, between nodes. Hardware networks, for example, may include clients, servers, and intermediary nodes in a graph topology (e.g., the user computer system 104, the merchant computer system 106, and the payment network system 108 and the transaction analysis system 110). In a similar fashion, data networks may include data nodes in a graph topology where each node includes related or linked information, software methods, and other data. A node such as the mobile computing device 104 may be managed to facilitate transactions with other nodes of the system (e.g., the merchant computer system) based on the data and instructions of various other nodes of the system 100.

(00026) It should be noted that the term “server” as used throughout this application refers generally to a computer, other device, program, or combination thereof that includes a processor and memory to process and respond to the requests of remote users/nodes across a communications network. Servers send their information to requesting “clients.” The term “client” as used herein refers generally to a computer,

program, other device, user and/or combination thereof that is capable of processing and making requests and obtaining and processing any responses from servers across a communications or data network. A computer, other device, set of related data, program, or combination thereof that facilitates, processes information and requests, and/or furthers the passage of information from a source user to a destination user is commonly referred to as a "node." Networks are generally thought to facilitate the transfer of information from source points to destinations. A node specifically tasked with furthering the passage of information from a source to a destination is commonly called a "router." There are many forms of networks such as Local Area Networks (LANs), Pico networks, Wide Area Networks (WANs), Wireless Networks (WLANs), etc. For example, the Internet is generally accepted as being an interconnection of a multitude of networks whereby remote clients and servers may access and interoperate with one another.

(00027) The mobile computing device 104 may include a smartphone 104 or other computing device capable of sending and receiving a wireless digital communication. The mobile computing device may include a processor 145, a memory 146, and a battery 147. The mobile computing device 104 may include an RFID tag 104A or other device capable of sending a radio signal including mobile computing device data such as a battery charge level 107. The memory 146 of the mobile computing device 104 may include various modules including instructions that, when executed by the processor 145 control the functions of the mobile computing device 104 generally and integrate the mobile computing device into the system 100 in particular. For example, some modules may include an operating system 150A, a browser module 150B, a communication module 150C, and a wallet module 150D. The communication module 150C may include processor-executable instructions to send and/or receive a signal from contactless component of the system. In some embodiments, the communication module 150C may include an RFID receiver or instructions to implement an RFID receiver. The wallet module 150D may also include payment device data 151. The payment device data 151 may securely identify a payment device 250 (Fig. 2) and a user's payment account information to facilitate a transaction between the mobile computing device 104 and the merchant computer system 106. In some embodiments,

the payment device data 151 may be tokenized such that only a trusted partner may access the payment device data 151.

(00028) The merchant computer system 106 may include a computing device such as a merchant server 129 including a processor 130 and memory 132 including components and instruction modules including processor-executable instructions to facilitate transactions with the mobile computing device 104 via other entities of the system 100. In some embodiments, the memory 132 may include a transaction communication module 134. The transaction communication module 134 may include instructions to send merchant messages 134A to other entities (i.e., 104, 108, 110) of the system 100 to indicate a transaction has been initiated with the mobile computing device 104 generally and wallet module 150D in particular including payment account data, location data, and other data as herein described. For example, a node of the system 100 (e.g., the mobile computing device 104 generally and the browser module 150B in particular) may access a merchant website to initiate a transaction. Some examples of transactions include ticket sales for specific seats at a venue.

(00029) The merchant computer system 106 may also include a transaction repository 142 and instructions to store payment and other transaction data 142A within the transaction repository 142. The merchant computer system 106 may also include a product repository 143 and instructions to store product and other data 143A within the product repository 143. In some embodiments, the merchant computer system 106 may also include instructions to send payment device data 151 corresponding to the payment device 250 (Fig. 2), transaction data 143A, and/or product data 143B and other data it received during a transaction to the payment network system from the mobile computing device 104.

(00030) The payment network system 108 may include a payment server 156 including a processor 158 and memory 160. The memory 160 may include a payment network module 162 including instructions to facilitate payment between parties (e.g., one or more users, merchants, etc.) using the system 100. The module 162 may be communicably connected to an account holder data repository 164 including payment network account data 164A. The payment network account data 164A may include any data to facilitate payment and other funds transfers between system entities (e.g., 104,

106 and 110). For example, the payment network account data 164A may include identification data, account history data, payment device data, etc. The module 162 may also include instructions to send payment messages 166 to other entities and components of the system 100 in order to complete transactions between the mobile computing system 104 and the merchant computer system 106. For example, the module 162 may include instructions to send a payment message 166 to a transaction analysis system 110 or other entity of the system 100 to complete a purchase transaction. The message 166 may include data to authorize a purchase transaction such as an authorization number or other identification, and may be tokenized or encrypted by the system 100 before the message 166 is sent over the network 102 to a system entity.

(00031) The transaction analysis system 110 may be used to provide real time or virtual real time analysis of a merchant's risk score. A server 116 may have a memory to store computer executable instructions which may physically configure a processor 118 according to the computer executable instructions. A table of transactions 112 may store individual transactions 112A, 112B, 112C and 112D for periodic or real time analysis. The transaction data may be stored in databases 112 and 124 which may contain past entries 122A and 124A which may be used to analyze the risk of present transactions.

(00032) Referring to Fig. 2, an embodiment of a system and method to create a real-time determination of a merchant risk level may be illustrated. As part of the determination, an entity 201 may be a user, a payment card, a merchant, a merchant category code (MCC), a merchant category group (MCG) or a postal code. Two types of entity profiles may be created in the current system – one based on domain knowledge, and the other based on a representation learning method.

(00033) DOMAIN KNOWLEDGE

(00034) Typical domain knowledge-based entity profiles may include historical fraud rate, decline rate, average dollar amount, etc. A typical example of such an entity profile - the distribution of fraud rate across different merchant category groups (MCGs) is shown below.

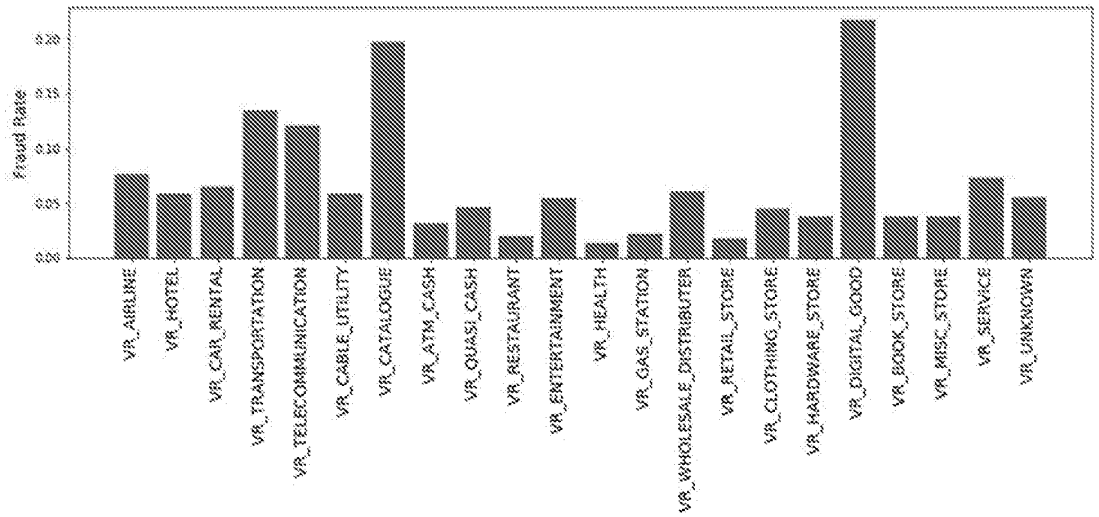

(00035) LEARNING PROFILE

(00036) Representation learning-based profiles may include user and merchant embeddings. An embedding may be a representation of a topological object, manifold, graph, field, etc. in a certain space in such a way that its connectivity or algebraic properties are preserved. User embeddings may be created by first collecting all the user-merchant pairs in transaction data and then grouping the pairs by distinct merchants. For each unique merchant, a list of all the transactions and the corresponding users may be found. The users may then be temporally ordered based on the time of transaction. Each user in the sequence may be treated as a "word" and the entire sequence may be treated as a "sentence". By repeating the same procedure across all the active merchants, a corpus of millions of sentences may be created. Word2vec or a similar algorithm may be used to generate an embedding for each user. Similarly, merchant embeddings may be created by first collecting all user-merchant pairs in transaction data and then grouping the pairs by distinct users. Each user embedding may be a 200-dimension float vector while each merchant embedding may be a 400-dimension float vector. The number of dimensions for each embedding representation may be chosen to be both efficient and effective for the task at hand. Both domain knowledge-based and representation learning-based entity profiles may be updated periodically.

(00037) Transaction Risk Model

(00038) The transaction risk model 202 may be trained offline in a supervised way using a combination of entity profiles and non-aggregated transaction specific features (such as location, time, dollar amount, etc.) as features and chargeback/fraud information as labels. The output of the model may be the probability of a transaction being fraudulent.

(00039) The probability may be converted into a risk score value between 1 and 100 where a higher score indicates a high-risk transaction, and a lower score indicates a low-risk transaction. To score each transaction, the proposed system may use a variety of models. Two models that may be used as an ensemble may include (i) Gradient Boosting Machine (GBM), and (ii) Recurrent Neural Network (RNN) including two long short-term memory (LSTM) layers.

(00040) At a high level, machine learning may be used to score transactions over time such that past scores may be analyzed in view of the resulting fraud to better determine the most appropriate scores in the future. In some embodiments, the scoring may be refined over time. Machine learning may be used to analyze past scores in view of the actual results of users or entities using credit. Machine learning may be used to review a training group of past scored data and determine weighting moving forward.

(00041) Fig. 6 may illustrate sample artificial intelligence (AI) training data according to one or more embodiments. As an example and not a limitation, an artificial intelligence system may be trained by analyzing a set of training data 605. The training data may be broken into sets, such as set A 610, set B 615, set C 620 and set D 625. As illustrated in Fig. 7a, one set may be used as a testing set (say set D 625) and the remaining sets may be used as training set (set A 610, set B 615 and set C 620). The artificial intelligence system may analyze the training set (set A 610, set B 615 and set C 620) and use the testing set (set D 625) to test the model created from the training data. Then the data sets may shift as illustrated in Fig. 7b, where the test data set may be added to the training data sets (say set A 610, set B 615 and set D 625) and one of the training data sets that have not been used to test before (say set C 620) may be used as the test data set. The analysis of the training data (set A 610, set B 615 and set D 625) may occur again with the new testing set (set C 620) being used to test the model and the model may be refined. The rotation of data sets may occur repeatedly

until all the data sets have been used as the test data sets. The model then may be considered complete and the model may then be used on additional data sets.

(00042) The RNN model architecture may be illustrated in Figure 4. LSTM 405 may process a cardholder's transaction sequence one at a time and may maintain a cell state vector $c(t)$ 410 and a hidden state vector $h(t)$ 415 which may act as a memory for past transaction behavior. The cell state vector $c(t)$ 410 may store information for longer periods of time, while the hidden state vector $h(t)$ 415 may store information for shorter periods of time. At each time step, LSTM 405 may reset, modify or mask the two state vectors 410 415 using in-built gating mechanisms. The two state vectors $c(t)$ 410 and $h(t)$ 415 at any time step t may be concatenated 420 into $\{c(t), h(t)\}$ which may then be updated based on previous state vectors $c(t-1)$, $h(t-1)$ and the current input vector $x(t)$:

(00043) $(c(t), h(t)) = LSTM\big(c(t-1), h(t-1), x(t)\big)$ $(Equation\ 1)$

(00044) In Equation 1, the current input vector $x(t)$ may consist of entity profiles created from the transaction logs.

(00045) The prediction from the deep recurrent neural network model may be obtained by

(00046) $\hat{y}(t) = softmax(Wh(t) + b)$ $(Equation\ 2)$

(00047) where W, b are learnable parameter arrays. The prediction $\hat{y}_t$ may act as a risk score indicating if a transaction is high-risk. The final risk score from a transaction may be a weighted average of the scores obtained from GBM and RNN models. The illustration below may present the performance curves of the GBM model, the RNN model, and the ensemble model.

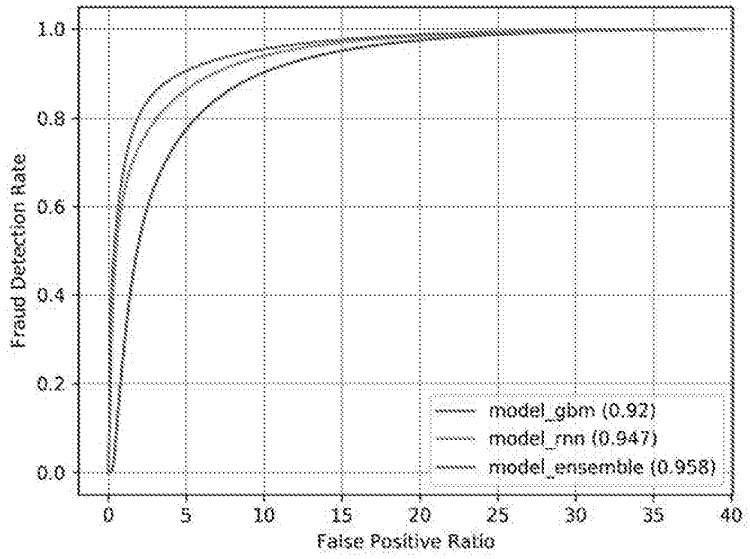

(00048) An ensemble of two models may result in noticeable performance improvement. The values in the parentheses may report the area-under-curve (AUC) score for the corresponding model. The higher the AUC value, the better the model performance.

(00049) Merchant Risk Model

(00050) The transaction risk model 202 introduced in the previous session may produce a risk score for each transaction that occurs at a merchant using a transaction score engine 206 (Fig. 2). The merchant risk model 203 may then aggregate the transaction risk scores corresponding to a merchant and may yield a merchant risk score in the range 1-100 using a merchant risk score engine 207 (Fig. 2). A higher score may indicate a high-risk merchant, and a lower score may indicate a low-risk merchant. Logically, the scores may be adjusted to reflect the desires of the user.

(00051) The system and method may attribute a risk score to a merchant in near real-time, but it may take ~ 1 day for an investigation team to review the merchants exhibiting anomalous behavior if needed. As illustrated in Fig. 5, the proposed system may aggregate risk scores on an hourly basis such as hour 1 (505), hour 2 (510) through hour 24 (515) by storing the hourly transactions in a bucket such as bucket 1 (525) for hour 1 (505), bucket 2 (530) for hour 2 (510) and bucket 24 (535) for hour 24

(515) and taking the mean of risk scores for transactions that occur within that hour. Then the system may compile the aggregated risk scores at the end of each day in a daily bucket (540) to generate a merchant risk score (550).

(00052) An anomaly detection algorithm 208 (Fig. 2) in an anomaly analysis model 204 may update the merchant risk score to generate a suspicious merchant list for review. The anomaly detection algorithm 208 may take on many forms. In one embodiment, the algorithm may simply look for merchant risk scores that are over a threshold. In another embodiment, the algorithm may highlight merchant risk scores which have changed more than expected over a time period, such as one day to the next. In yet another embodiment, the algorithm may highlight merchant risk scores which have experienced the greatest percentage change over a period of time. Of course, all or some of the different measurements described may be part of a single algorithm and may be given different weights to determine which merchants may need to be studied more closely by a merchant case manager 209.

(00053) Anomaly Analysis Model 204

(00054) One contributions of the proposed solution may be the near-real time anomaly detection capability for millions of merchants in a payment network. The daily aggregated merchant risk scores may be viewed as a time series, as shown below.

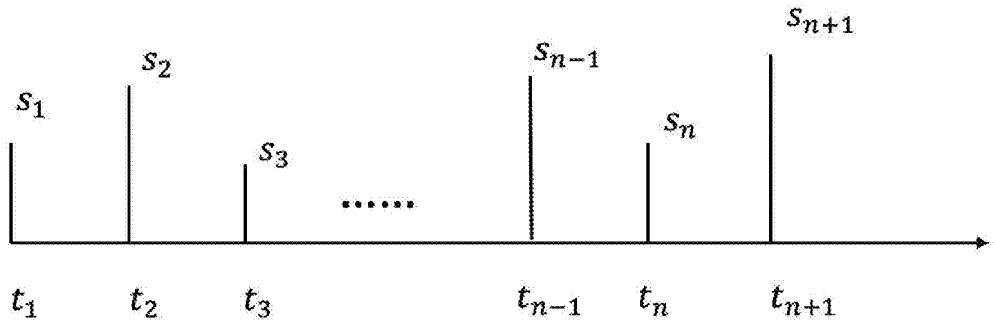

(00055) For the daily aggregated merchant risk scores, the anomaly analysis model 204 may utilize a recursive exponential filter to identify anomalous merchant behavior. The recursive exponential filter, which may be analogous to the widely known Kalman filter, may eliminate the need to store and reprocess past transaction records each time the risk score needs to be updated. Below may be an illustration of a time series

showing a sequence of pairs $\{s_1, t_1\}, \{s_2, t_2\}, \ldots, \{s_n, t_n\}, \{s_{n+1}, t_{n+1}\}$, where s_i is the merchant risk score at time t_i.

(00056) Let $\{s_i, t_i\}$ represent a sequence of pairs, where s_i is the daily aggregated merchant risk score for day t_i. The merchant risk level on day t_n may be defined as a running average y_n which may then be given by

(00057) $$y_n = \sum_{i=1}^{n} e^{-(t_n - t_i)/T} s_i / \sum_{i=1}^{n} e^{-(t_n - t_i)/T} \qquad (Equation\ 3)$$

(00058) where T is a decay parameter that controls how quickly should the model may forget about the past history. After the next daily merchant risk score observation at time t_{n+1}, the merchant risk level y_{n+1}may be calculated as

(00059) $$y_{n+1} = \sum_{i=1}^{n+1} e^{-(t_{n+1} - t_i)/T} s_i / \sum_{i=1}^{n+1} e^{-(t_{n+1} - t_i)/T} \qquad (Equation\ 4)$$

(00060) With slight manipulation of Equation 3 and Equation 4, the system and method may represent y_{n+1} using the following recursive formula:

(00061) $$y_{n+1} = y_n + K_{n+1}(s_{n+1} - y_n) \qquad (Equation\ 5)$$

(00062) where K_{n+1} may also be cast in a recursive formula:

(00063) $$K_{n+1} = K_n / \left(K_n + e^{-(t_{n+1} - t_n)/T}\right) \qquad (Equation\ 6)$$

(00064) A significant deviation of the merchant risk level y_{n+1} from y_n may indicate an anomaly. It is important to notice an entire history of merchant risk scores may not need to be stored. Instead, the previous risk score s_n, previous merchant risk level y_n, and K_n may only be needed to be stored.

(00065) Figure 8, 9 and 10 present three example of merchants exhibiting anomalous behavior. The first example (Fig. 8) may illustrate a point anomaly where a single spike 805 may be seen in the graph.

(00066) The second example (Fig. 9) may illustrate a score shift 905 (from low to high) followed by two spikes.

(00067) The third example (Figure 10) may illustrate a sudden jump in risk score 1005 and the score stays high for a period of time.

(00068) How to evaluate new Merchants?

(00069) In production, "peer group" risk scores may be maintained. A peer group may be a collection of merchants that provide similar goods or services. For instance, all restaurants may belong to the same peer group. There may be a one-to-one mapping between a merchant category code (MCC) and a peer group. There may be ~750 MCC codes in traditional transaction data, which may correspond to ~750 peer groups. The ~750 MCC groups may be further classified into 22 merchant category groups (MCGs) in the system. For instance, the restaurant MCC (5812) and fast-food restaurant MCC (5814) may be assigned the same MCG (10). The hierarchical classification of merchants at different granularities ensures that peer groups may be produced with robust statistics. If a merchant category code (MCC) has a small number of transactions, the system and method may move up one layer to use the merchant category group (MCG). Similar to calculating the merchant risk score, all transaction risk scores in the peer group may be aggregated on an hourly basis and the aggregated risk scores may be compiled at the end of each day to produce a peer group risk score. Peer group analysis may help assess the risk of new merchants that have short historical data.

(00070) A computer implemented method for real-time determination of a merchant risk level may be illustrated in FIG. 3. At block 200, a transaction risk score may be calculated for each transaction that occurs at a merchant during an hour h, wherein h is from 1 to 24. The transaction risk model 202 may be trained offline in a supervised way using a combination of entity profiles and non-aggregated transaction specific features (such as location, time, dollar amount, etc.) as features and chargeback/fraud information as labels. The output of the model may be the probability of a transaction being fraudulent.

(00071) The probability may be converted into a risk score value between 1 and 100 where a higher score indicates a high-risk transaction, and a lower score indicates a

low-risk transaction. To score each transaction, the proposed system may use a variety of models. Two models that may be used as an ensemble may include (i) Gradient Boosting Machine (GBM), and (ii) Recurrent Neural Network (RNN) including two long short-term memory (LSTM) layers.

(00072) At block 210 the transaction risk scores for each transaction that occurred at the merchant during hour h may be aggregated to obtain a merchant risk score m_h for the merchant during each hour h as illustrated in Fig. 5. At block 220, the calculating step and the aggregating step may be repeated for each hour h.

(00073) At block 230, the merchant risk scores m_h (535, 530, 525) may be compiled (540) to calculate a daily merchant risk score s_i (550). The merchant risk model 203 may then aggregate the transaction risk scores corresponding to a merchant and may yield a merchant risk score in the range 1-100 using a merchant risk score engine 207 (Fig. 2). A higher score may indicate a high-risk merchant, and a lower score may indicate a low-risk merchant. Logically, the scores may be adjusted to reflect the desires of the user. In some embodiments, a mean of the transaction risk scores may be calculated for all the transactions that occurred during hour h.

(00074) At block 240 the daily merchant risk score s_i (550) may be provided to the merchant. In response, the merchant may see virtually immediately when the merchant risk score is approaching an alert level. In some embodiments, the system and method may identify the source of a change in the merchant risk score. For example, if a certain card is causing problems, the merchant may make an effort to look for that card. In other embodiments, purchases during a certain period of time may be a problem and the merchant may take additional risk reduction steps during those periods of time. In yet another embodiment, the merchant may be informed the sale of certain items has been identified as risky and the merchant may adjust the sales of those good.

(00075) As can be seen, the system and method may address the technical problem of how to check a merchant rating score in a computationally efficient manner. In the past, checking a merchant rating in a timely manner was not even possible. Theoretically, a user may have made a daily request for an update which may or may not have resulted in a response as providing daily updates is simply not done at the present time. Logically, merchant aggregation is also not done at the present time. At a

minimum, the system and method enable some new functionality which was not possible previously, using any computer system.

(00076) In some further embodiments, the method may also determine a peer group for the merchant, where the peer group includes one or more peer merchants that provide similar goods or services as the merchant. A transaction risk score may be calculated for each transaction that occurs at the peer group during an hour h, wherein h is from 1 to 24. The transaction risk score for each transaction that occurred at the peer group during hour h may be aggregated to obtain a peer group risk score p_h for the peer group. The calculating step and the aggregating step may be repeated for each hour h. The peer group scores p_h may be compiled to calculate a daily peer group risk score p_i. The daily peer group risk score p_i may be compared to the daily merchant risk score s_i to better determine the level of transaction risk for a single merchant in view of the world of similar merchants.

(00077) In some embodiments, an anomaly detection algorithm may be used to determine if an anomaly exists in the daily merchant risk score s_i. If an anomaly is detected, the merchant may be added to a list. The anomaly detection algorithm may include determining a first merchant risk level at a time t_n, wherein

(00078) the first merchant risk level y_n is a running average of the merchant daily risk scores si, wherein the first merchant risk level y_n is calculated as,

(00079) $$y_n = \sum_{i=1}^{n} e^{-(t_n - t_i)/T} s_i / \sum_{i=1}^{n} e^{-(t_n - t_i)/T}$$

(00080) And wherein T is a decay parameter and i is from 1 to n.

(00081) A second merchant risk level at a time t_{n+1} may be determined wherein the second merchant risk level y_{n+1} is calculated as,

(00082) $$y_{n+1} = \sum_{i=1}^{n+1} e^{-(t_{n+1} - t_i)/T} s_i / \sum_{i=1}^{n+1} e^{-(t_{n+1} - t_i)/T}$$

(00083) The first merchant risk level y_n may be compared to the second merchant risk level y_{n+1}. If the difference between the first merchant risk level y_n and the second merchant risk level y_{n+1} is above a threshold, then an anomaly may be identified.

(00084) As mentioned previously, machine learning may be used to determine if a transaction is likely fraudulent. A first entity profile may be created using a first dataset. A second entity profile may be created using a second dataset, where the second entity profile is created using the first entity profile and the second entity profile may be used to train a third dataset containing transaction data to determine a probability that a transaction is fraudulent. The probability that the transaction is fraudulent may be converted into a transaction risk score. The second entity profile may include user and merchant embeddings and the first entity profile may include one or more of a historical fraud rate, a transaction decline rate, and an average dollar amount. The transaction risk score may be any of 1 to 100 for example. Logically, a high transaction risk score may indicate a high-risk transaction and a low transaction risk score may indicate a low risk transaction.

(00085) The system 400 may include but is not limited to any combination of a LAN, a MAN, a WAN, a mobile, a wired or wireless network, a private network, or a virtual private network. Moreover, while only one remote computing device 104 is illustrated in Fig. 1 to simplify and clarify the description, it is understood that any number of client computers are supported and can be in communication within the system 100.

(00086) Additionally, certain embodiments are described herein as including logic or a number of components, modules, blocks, or mechanisms. Modules and method blocks may constitute either software modules (e.g., code or instructions embodied on a machine-readable medium or in a transmission signal, wherein the code is executed by a processor) or hardware modules. A hardware module is tangible unit capable of performing certain operations and may be configured or arranged in a certain manner. In example embodiments, one or more computer systems (e.g., a standalone, client or server computer system) or one or more hardware modules of a computer system (e.g., a processor or a group of processors) may be configured by software (e.g., an application or application portion) as a hardware module that operates to perform certain operations as described herein.

(00087) In various embodiments, a hardware module may be implemented mechanically or electronically. For example, a hardware module may comprise dedicated circuitry or logic that is permanently configured (e.g., as a special-purpose

processor, such as a field programmable gate array (FPGA) or an application-specific integrated circuit (ASIC)) to perform certain operations. A hardware module may also comprise programmable logic or circuitry (e.g., as encompassed within a processor or other programmable processor) that is temporarily configured by software to perform certain operations. It will be appreciated that the decision to implement a hardware module mechanically, in dedicated and permanently configured circuitry, or in temporarily configured circuitry (e.g., configured by software) may be driven by cost and time considerations.

(00088) Accordingly, the term "hardware module" should be understood to encompass a tangible entity, be that an entity that is physically constructed, permanently configured (e.g., hardwired), or temporarily configured (e.g., programmed) to operate in a certain manner or to perform certain operations described herein. As used herein, "hardware-implemented module" refers to a hardware module. Considering embodiments in which hardware modules are temporarily configured (e.g., programmed), each of the hardware modules need not be configured or instantiated at any one instance in time. For example, where the hardware modules comprise a processor configured using software, the processor may be configured as respective different hardware modules at different times. Software may accordingly configure a processor, for example, to constitute a particular hardware module at one instance of time and to constitute a different hardware module at a different instance of time.

(00089) Hardware modules may provide information to, and receive information from, other hardware modules. Accordingly, the described hardware modules may be regarded as being communicatively coupled. Where multiple of such hardware modules exist contemporaneously, communications may be achieved through signal transmission (e.g., over appropriate circuits and buses) that connect the hardware modules. In embodiments in which multiple hardware modules are configured or instantiated at different times, communications between such hardware modules may be achieved, for example, through the storage and retrieval of information in memory structures to which the multiple hardware modules have access. For example, one hardware module may perform an operation and store the output of that operation in a memory device to which it is communicatively coupled. A further hardware module may

then, at a later time, access the memory device to retrieve and process the stored output. Hardware modules may also initiate communications with input or output devices, and can operate on a resource (e.g., a collection of information).

(00090) The various operations of example methods described herein may be performed, at least partially, by one or more processors that are temporarily configured (e.g., by software) or permanently configured to perform the relevant operations. Whether temporarily or permanently configured, such processors may constitute processor-implemented modules that operate to perform one or more operations or functions. The modules referred to herein may, in some example embodiments, comprise processor-implemented modules.

(00091) Similarly, the methods or routines described herein may be at least partially processor-implemented. For example, at least some of the operations of a method may be performed by one or processors or processor-implemented hardware modules. The performance of certain of the operations may be distributed among the one or more processors, not only residing within a single machine, but deployed across a number of machines. In some example embodiments, the processor or processors may be located in a single location (e.g., within a home environment, an office environment or as a server farm), while in other embodiments the processors may be distributed across a number of locations.

(00092) The one or more processors may also operate to support performance of the relevant operations in a "cloud computing" environment or as a "software as a service" (SaaS). For example, at least some of the operations may be performed by a group of computers (as examples of machines including processors), these operations being accessible via a network (e.g., the Internet) and via one or more appropriate interfaces (e.g., application program interfaces (APIs).)

(00093) The performance of certain of the operations may be distributed among the one or more processors, not only residing within a single machine, but deployed across a number of machines. In some example embodiments, the one or more processors or processor-implemented modules may be located in a single geographic location (e.g., within a home environment, an office environment, or a server farm). In other example

embodiments, the one or more processors or processor-implemented modules may be distributed across a number of geographic locations.

(00094) Some portions of this specification are presented in terms of algorithms or symbolic representations of operations on data stored as bits or binary digital signals within a machine memory (e.g., a computer memory). These algorithms or symbolic representations are examples of techniques used by those of ordinary skill in the data processing arts to convey the substance of their work to others skilled in the art. As used herein, an "algorithm" is a self-consistent sequence of operations or similar processing leading to a desired result. In this context, algorithms and operations involve physical manipulation of physical quantities. Typically, but not necessarily, such quantities may take the form of electrical, magnetic, or optical signals capable of being stored, accessed, transferred, combined, compared, or otherwise manipulated by a machine. It is convenient at times, principally for reasons of common usage, to refer to such signals using words such as "data," "content," "bits," "values," "elements," "symbols," "characters," "terms," "numbers," "numerals," or the like. These words, however, are merely convenient labels and are to be associated with appropriate physical quantities.

(00095) Unless specifically stated otherwise, discussions herein using words such as "processing," "computing," "calculating," "determining," "presenting," "displaying," or the like may refer to actions or processes of a machine (e.g., a computer) that manipulates or transforms data represented as physical (e.g., electronic, magnetic, or optical) quantities within one or more memories (e.g., volatile memory, non-volatile memory, or a combination thereof), registers, or other machine components that receive, store, transmit, or display information.

(00096) As used herein any reference to "embodiments," "some embodiments" or "an embodiment" or "teaching" means that a particular element, feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment. The appearances of the phrase "in some embodiments" or "teachings" in various places in the specification are not necessarily all referring to the same embodiment.

(00097) Some embodiments may be described using the expression "coupled" and "connected" along with their derivatives. For example, some embodiments may be described using the term "coupled" to indicate that two or more elements are in direct physical or electrical contact. The term "coupled," however, may also mean that two or more elements are not in direct contact with each other, but yet still co-operate or interact with each other. The embodiments are not limited in this context.

(00098) Further, the figures depict preferred embodiments for purposes of illustration only. One skilled in the art will readily recognize from the following discussion that alternative embodiments of the structures and methods illustrated herein may be employed without departing from the principles described herein

(00099) Upon reading this disclosure, those of skill in the art will appreciate still additional alternative structural and functional designs for the systems and methods described herein through the disclosed principles herein. Thus, while particular embodiments and applications have been illustrated and described, it is to be understood that the disclosed embodiments are not limited to the precise construction and components disclosed herein. Various modifications, changes and variations, which will be apparent to those skilled in the art, may be made in the arrangement, operation and details of the systems and methods disclosed herein without departing from the spirit and scope defined in any appended claims.