WILEY | Hindawi

Research Article

Credit Card Fraud Detection through Parenclitic Network Analysis

Massimiliano Zanin,[1,2,3] Miguel Romance [iD],[3,4,5] Santiago Moral,[3,4,6] and Regino Criado[3,4,5]

[1]*Department of Computer Science, Faculty of Science and Technology, Universidade Nova de Lisboa, Lisboa, Portugal*
[2]*Center for Biomedical Technology, Universidad Politécnica de Madrid, 28223 Pozuelo de Alarcón, Madrid, Spain*
[3]*Data, Networks and Cybersecurity Research Institute, Univ. Rey Juan Carlos, 28028 Madrid, Spain*
[4]*Department of Applied Mathematics, Universidad Rey Juan Carlos, 28933 Móstoles, Madrid, Spain*
[5]*Center for Computational Simulation, 28223 Pozuelo de Alarcón, Madrid, Spain*
[6]*Cyber Security & Digital Trust, BBVA Group, 28050 Madrid, Spain*

Correspondence should be addressed to Miguel Romance; miguel.romance@urjc.es

The detection of frauds in credit card transactions is a major topic in financial research, of profound economic implications. While this has hitherto been tackled through data analysis techniques, the resemblances between this and other problems, like the design of recommendation systems and of diagnostic/prognostic medical tools, suggest that a complex network approach may yield important benefits. In this paper we present a first hybrid data mining/complex network classification algorithm, able to detect illegal instances in a real card transaction data set. It is based on a recently proposed network reconstruction algorithm that allows creating representations of the deviation of one instance from a reference group. We show how the inclusion of features extracted from the network data representation improves the score obtained by a standard, neural network-based classification algorithm and additionally how this combined approach can outperform a commercial fraud detection system in specific operation niches. Beyond these specific results, this contribution represents a new example on how complex networks and data mining can be integrated as complementary tools, with the former providing a view to data beyond the capabilities of the latter.

1. Introduction

Credit card fraud, a concept included in the wider notion of financial frauds [1, 2], is a topic attracting an increasing attention from the scientific community. This is due, on the one hand, to the rising costs that they generate for the system, reaching billions of dollars in yearly losses and a percentage loss of revenues equal to the 1.4% of on-line payments [3]. On the other hand, credit card frauds have important social consequences and ramifications, as they support organized crime, terrorism funding, and international narcotics trafficking; see [4] for a complete review.

Detecting unauthorized credit card transactions is an extremely complex problem, as features are seldom useful if taken individually. To illustrate, a large transaction is not *prima facie* suspicious, unless it is performed at usual times

(e.g., at night) or in an unusual store (i.e., a store never visited before by the card owner, located in a different city, etc.). When different features have to be combined in non-trivial ways, the customary solution is to resort to data mining, a subfield of computer science dealing with the automatic discovery of patterns in data sets [5–7]. While data mining algorithms are able to detect hidden patterns in data, they usually lack the capacity of synthesizing metrics describing the global structure created by the interactions between the different features. In recent years, the use of complex networks theory has been proposed as a way of overcoming this limitation. Complex networks are a statistical-mechanics understanding of the classical graph theory, aimed at describing and characterizing the structure of complex systems [8–10]. The interaction between network theory and data mining is bidirectional: the former can be used to synthesize

high-level features to be fed into a classification problem, while the latter can endow networks with an objective way of validating results; see [11] for a complete review.

More specifically, complex networks and data mining can be integrated as complementary tools in order to extract, synthesize, and create new representations of a data source, with the aim of, for instance, discover new hidden patterns in a complex structure. The appropriate integration of complex network metrics can result in improved classification rates with respect to classical data mining algorithms and, reciprocally, there are many situations in which data mining can be used to solve important issues in complex network theory and applications [11].

In this contribution we explore the possibility of using complex networks as a way of improving credit card fraud detection. Specifically, networks are used to synthesize complex features representing card transactions, relying on the recently proposed approach of *parenclitic networks* (Section 3). Afterwards, their relevance is evaluated by means of a large dataset of real transactions, by comparing the yielded increase in the classification score when compared to the use of a standard ANN algorithm (Section 4). We additionally show that the combined data mining/complex networks approach is able to outperform a commercial system in some specific situations.

2. State of the Art in Credit Card Frauds Detection

Due to the high importance, both economic and social, of the problem of detecting frauds in credit card transaction, it is not surprising that a large body of works can be found in the literature, especially based on the analysis of past transaction data. While a complete review is out of the scope of this contribution, in this section we review some important techniques; the interested reader may refer to [32–34] for more comprehensive reviews.

Generally speaking, credit card frauds detection approaches can be classified in two main families, which correspond to the two families of machine learning algorithms: supervised and unsupervised ones. In the former family, past transactions are labeled as legal or illegal, for instance, based on expert judgement or customer's claims; the algorithms then learn over these data, to create a model that is applied to new instances appearing in the system. On the other hand, unsupervised techniques are based on the automatic detection of patterns that are considered "normal" for a given user, for then detecting transactions that are not coherent with such patterns, as illegal users are expected to depart from the owner's behaviors. Both families have their own advantages and disadvantages, and in general supervised approaches are more effective in detecting illegal transactions, although they require a large initial training set (which in some cases may not be available).

Table 1 reports some relevant works, organized by the machine learning technique used. Most of the data mining models to detect credit card frauds are based on artificial neural networks (ANNs), a model inspired on the structural

Table 1: List of relevant references in credit card frauds detection, based on data analysis and machine learning, organised according to the algorithm used.

Algorithm	Type	References
Artificial Neural Networks	Supervised	[12–18]
Self-Organising Maps	Unsupervised	[19, 20]
Genetic Algorithm	Supervised	[21–23]
Support Vector Machines	Supervised	[24–27]
Bayesian networks	Supervised	[15]

aspects of biological neural networks, and in which a set of nodes process the input signal by interacting between them [35, 36]. The preference of this algorithm is based on the fact that ANNs are able to extract complex nonlinear patterns from data, with almost no hypotheses on the underlying structure; they are thus a natural choice, albeit with some limitations, including a high computational cost, and the fact of being "black-boxes". Other relevant supervised algorithms include genetic algorithms (GA), in which a set (or a *population*) of solutions are evolved using rules inspired on genetic natural selection [37, 38]; support vector machines (SVMs), a classification algorithm based on finding the hyper-plane in the feature hyper-space able to divide instances according to their classes [39, 40]; and Bayesian networks, probabilistic models representing relationships between features and classes through directed acyclic graphs [41, 42]. As for unsupervised algorithms, it is worth citing self-organizing maps (SOM), a type of ANN whose output is tuned to be a low-dimensional representation of the input features [43].

In spite of the large number of publications focusing on the problem of detecting illegal transactions, the research community still faces some important problems. First, there are no public and creditable data sets against which algorithms can be tested and benchmarked. This, of course, is a major obstacle towards reaching high levels of reproducibility, but, on the other hand, the privacy concerns about credit card data are a barrier difficult to overcome; see [44] for a discussion. Second, there is no accepted and common way of measuring the effectiveness of a classification model, with previous works heterogeneously relying on accuracy, precision, area under the ROC curve [45], or F1-measures.

3. Methods

In this section we present the main tools that are going to be used for the classification of credit card transactions between licit and illicit. Given a credit card transaction t_i with features $f_{i1}, \ldots, f_{ik}$, the problem entails detecting if it is illicit or not from its features and the knowledge obtained from a historical training dataset, what is known as a supervised learning problem. From a mathematical point of view, we have to model a function $H : \mathbb{R}^k \to \mathbb{R}$ and find $\delta > 0$ such that if $|H(f_{i1}, \ldots, f_{ik})| \leq \delta$, then t_i is not illicit. Note that while there are multiple types of illicit patterns, such aspect is here not considered, in that any suspicious transaction is considered as a potential fraudulent one.

We firstly introduce the concept of *parenclitic networks* in Section 3.1, a network reconstruction technique that allows highlighting the differences between one instance and a set of standard (i.e., baseline, or in this case licit) instances [46, 47]. We subsequently describe the real data set used for validation (Section 3.2), including the available raw features (Table 3); and the global classification model (Section 3.3).

3.1. Parenclitic Networks Reconstruction. As initially proposed in [46], one may hypothesise that the right classification of an observation does not only come from its features, but also comes from the structure of correlations between them. Following the mathematical formalism introduced before, if we consider the set

$$L = \left\{ (x_1, \ldots, x_k) \in \mathbb{R}^k;\ |H(x_1, \ldots, x_k)| \le \delta \right\} \subseteq \mathbb{R}^k, \quad (1)$$

then L is a manifold in $\mathbb{R}^k$ such that if we take a (new) transaction t with features $t_1, \ldots, t_k$ such that $(t_1, \ldots, t_k) \notin L$; then t is considered as an illicit transaction. In general it is computationally impossible to obtain the set L directly from the training dataset, since it is a high dimensional problem. As an alternative, the parenclitic approach analyzes the family of projections of L into 2-dimensional spaces corresponding to couples of features (x_i, x_j) with $1 \le i \ne j \le k$. Hence, if we consider a training dataset with $n \in \mathbb{N}$ transactions, each of them described by $k \in \mathbb{N}$ (numeric) features, we can analyze up to $\binom{k}{2} = k(k-1)/2$ two-dimensional projections of pairs of different features, each of them with up to n points in $\mathbb{R}^2$. In order to quantify the correlation between pairs of features, the parenclitic approach proposes associating a network to each transaction with k nodes (as many as features considered) and the links measure the correlation between features [47]. Hence the following preprocessing must be completed: for every two-dimensional projection of L given by a couple of features (f_i, f_j) with $1 \le i \ne j \le k$, the correlation for the licit transactions in the training dataset is measured (by means of, for instance, a linear regression or other curve fitting techniques). For the sake of simplicity, we have here considered a linear regression, such that every pair of features (f_i, f_j) with $1 \le i \ne j \le k$ yields a linear fitting between f_i and f_j for the licit transactions in the training dataset. Mathematically, this is represented by a linear equation of the form

$$r_{ij} : x_j = a_{ij} x_i + b_{ij}. \quad (2)$$

Once these $\binom{k}{2}$ linear regression lines are computed, a threshold $\alpha > 0$ is fixed. Given a new (i.e., not included in the training set) transaction t with features $t_1, \ldots, t_k$, a network $G = G(t)$ is associated with t as follows:

(i) G has k nodes $1, \ldots, k$,

(ii) For every pair of nodes $1 \le i \ne j \le k$ we compute $w_{ij} \ge 0$ as the (Euclidean) distance from (t_i, t_j) to the line r_{ij} in $\mathbb{R}^2$, i.e.,

$$w_{ij} = d\left((t_i, t_j), r_{ij} \right). \quad (3)$$

As an alternative, the Euclidian distance could be replaced by any pseudo-distance function in $\mathbb{R}^2$. For the sake of simplicity, the Euclidian distance will be used in this paper, but similar results can be obtained for other pseudo-distance functions.

(iii) For every pair of nodes $1 \le i \ne j \le k$, the (undirected) link (i, j) is in graph G if and only if $w_{ij} \ge \alpha$.

Note that the parenclitic network $G(t)$ summarizes the couples of features whose correlation strongly differs from a typical licit transaction; the structure of this network thus contains valuable information about the (abnormal) correlation of features in the credit card transaction. Once this parenclitic network is computed, it is necessary to transform it in a set of features compatible with a data mining algorithm. Towards this end, several structural measures have been extracted and will be considered as new features associated with the transaction (see next section for details).

The identification of the best set of measures to describe the network representation of a system is not a trivial nor a closed problem. Researchers usually resort to two different strategies: define a set of metrics, according to the structural aspects they describe and to the known properties of the system under study or relying in an external optimization phase, for instance, by including a large set of metrics, and by applying a feature selection algorithm on them [11, 48]. In this contribution we choose the first option, as (i) it minimizes the probability of overfitting, and (ii) the expected structure of the parenclitic networks has already been studied. Regarding the last point, it has been shown that the networks corresponding to abnormal (in the sense of different from the average) instances have a star-like structure, usually characterized by high clustering coefficient (number of triangles), high efficiency, and low Information Content [46, 47]. By using this initial information, a set of relevant features have been selected among all possible structural measures that could be computed (see, for example, [28] and references therein), as summarized in Table 2.

3.2. Data Set Description. The data set here considered includes all credit and debit card transactions of clients of the Spanish bank BBVA, from January 2011 to December 2012. Each month, an average of 15 million operations were realized by 7 million cards, for a total of 250 GB of information.

Transactions are automatically screened by an algorithm designed to detect suspected transactions, and returning a score from 0 (no suspect) to 100 (potentially illegal). Afterwards, transactions are classified in two categories, i.e., *legal* and *illegal*, as a result of a manual classification performed by the bank's legal personnel, using both information of the automatic algorithm and customers' complaints. This allows us to detect which transactions were positively detected as frauds by the automatic algorithm, and which were false negatives.

Available fields included a time stamp of the operation, the quantity (both in Euro and in the original currency, if different), and the origin (the card) and destination (the store) of the operation; the two latter fields were anonymized, so that

TABLE 2: List of topological metrics used to describe the structure of parenclitic networks.

Name	Description
Maximum node degree [28]	Maximum degree of all nodes in the network. It is calculated as $M_k = \max_i k_i$, k_i being the degree of nodes i
Entropy of the degree distribution [29]	Shannon entropy of the distribution of nodes degrees. It is given by $E = -\sum_{i=0}^{M_k} p_i \log p_i$, p_i being the probability of finding a node of degree i.
Assortativity [28]	Pearson's correlation coefficient between the degree of connected nodes.
Clustering coefficient [28]	Measure of the presence of triangles in the network. It is defined as the number of triangles (groups of three fully-connected nodes) over the number of connected triplets (groups of three nodes connected by at least two links).
Geodesic distance [28]	Average length of the shortest path connecting pairs of nodes.
Efficiency [30]	Inverse of the harmonic mean of the length of all shortest distances.
Information Content [31]	Metric assessing the presence of meso-scale structures in the network.

TABLE 3: Features composing the credit card transactions dataset.

Name	Type	Description
Transaction size	Integer	Size, in Euro, of the transaction under analysis.
Time since last transaction	Integer	Time, in seconds, since the last transaction of the same card.
Last transaction size	Integer	Size, in Euro, of the previous transaction executed by the same card.
Average transaction size	Float	Average size, in Euro, of the transactions executed by the card in the last month.
Average time between transactions	Float	Average time, in seconds, between consecutive transactions of the same card.
Same shop	Boolean	1 is the shop corresponds to the one of the last transaction of the same card, 0 otherwise.
Hour of the day	Integer	Hour (from 1 to 24) at which the operation was realized.
Fraud rate	Float	Average rate of illegal operations, for all cards, in the last 50.000 transactions.
Fraud suspicion	Integer	Number representing the likelihood for the transaction to be illicit, according to the bank automatic fraud detection algorithm. Values range between 0 (no fraud suspected) to 100 (certain fraud).
Fraud	Boolean	1 if the transaction has been recognized as a fraud, 0 otherwise.

the exact card number and the name of the store could not be recovered. Some additional features have been synthesized from the previous ones, e.g., the average transaction size of a given user. A full list of the available fields is reported in Table 3. Additionally, a full statistical characterization of the features can be found in [49], including the temporal evolution of the structure of the transactions network.

3.3. Classification Models.

As previously introduced, in this contribution we are going to explore two different ways of detecting illicit credit card transactions: a classical data mining approach, and the introduction of features extracted from a network representation. In both cases, the process must follow some common steps: it is first necessary to extract the expected behavior, i.e., a set of features representing the typical legal and illegal transaction; by using such features, a mathematical model will be constructed in order to learn the differences between legal and illegal transactions. Once this mathematical model is fitted, new transactions not studied previously could be classified.

Figure 1 depicts an overview of the whole process. It starts from the original data set, from which a set of raw features are extracted, described in Section 3.2 and listed in Table 3. The features corresponding to the licit transactions are then used
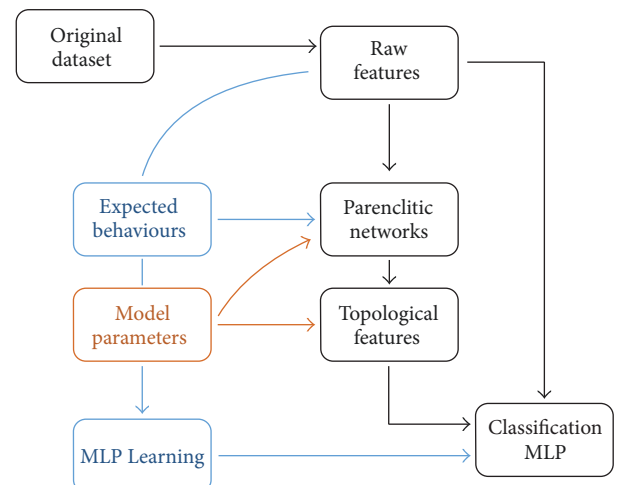

FIGURE 1: Schematic representation of the classification model. See main text for details.

to recover the normal relations, as described in Section 3.1, and to reconstruct the parenclitic networks of all transactions. These networks are then binarized; i.e., links with weight
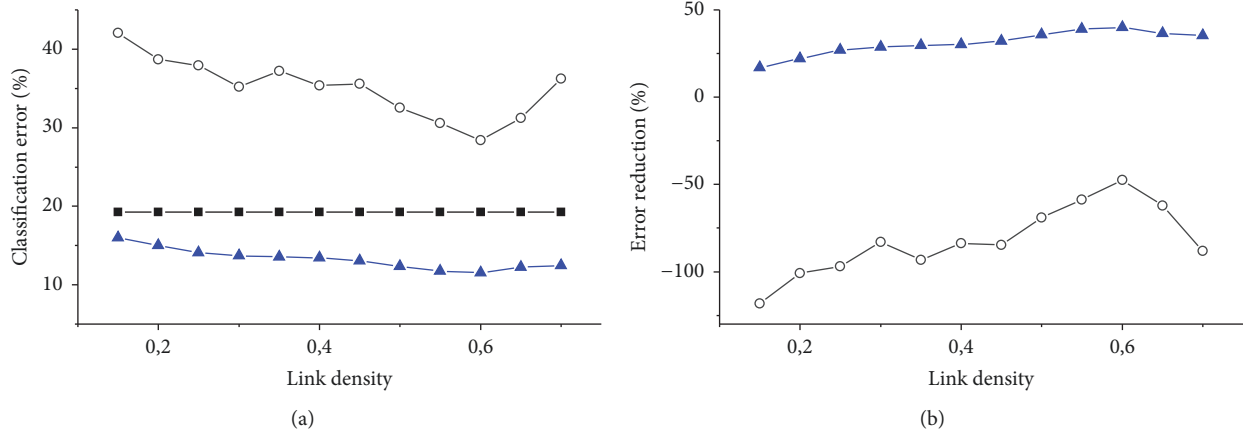
FIGURE 2: (a) Classification error as a function of the link density of the parenclitic networks. Black squares, black circles, and blue triangles, respectively, represent the error for the classification for the original raw features, for the classification using parenclitic features alone, and for the classification with all features. (b) Error reduction, in percentage, when using only parenclitic features (black circles) and the full set of features (blue triangles), with respect to the use of the raw data set.

below a given threshold are deleted, and a set of topological metrics are extracted; see Table 2 for a complete list. Note that, at the end of this analysis, all transactions are described by 15 features: 8 coming from the raw data, and 7 from the network analysis.

Artificial Neural Networks (ANNs), and specifically Multi-Layer Perceptrons (MLP), have been chosen as the final model for classifying new transactions. They are inspired by the structural aspects of biological neural networks and are represented by a set of connected nodes in which each connection has a weight associated with it, and the network learns the classification function adjusting the node weights [36, 50]. The output of each artificial neuron j is defined by

$$f\left(W^T, x_j\right) = \sum_{i=1}^{n} W_i x_i + b, \tag{4}$$

W being the vector of weights and f the sigmoid activation function:

$$f(x) = \frac{1}{1 + \exp(-x)}. \tag{5}$$

Following the standard configuration, neurons were organized in three layers: an input one, with a number of neurons equal to the input features; an intermediate, or hidden one, with ten neurons; and a final output layer comprising just one computational element. The training has been performed with the standard back-propagation algorithm [51]. Finally, the reconstruction of the MLP models has been performed using the KNIME software [52].

The evaluation of the classification efficiency has been performed using both sensitivity (also known as True Positive Rate (TPR)) and Receiver Operating Characteristic (ROC) curves [45]. These curves are created by plotting the True Positive Rate (TPR) against the False Positive Rate (FPR) at various threshold settings. ROC plots present the important advantage of showing the performance of the classification model for different sensitivity values. This is relevant for the

problem at hand, as false positives are extremely expensive, e.g., in terms of the negative commercial image of the bank; conservative solutions are therefore usually preferred. All results here presented correspond to a cross-validation, in which data of the first year (approximatively half of the instances) are used for training, and those of the second year for evaluation only. This, together with the high number of available instances, minimizes the risk of model overfitting [53].

4. Results

As explained in Section 3.1, the parenclitic approach usually requires the definition of a threshold α, which is used to binarize the (initially weighted) networks. Instead of using an a priori approach, i.e., the definition of α using expert judgement, we here tackle the problem indirectly, by following the procedure proposed in [48]. Specifically, we optimize the network reconstruction by finding the link density (and hence the value of α) that optimizes the efficacy of the classification model.

Figure 2(a) presents the evolution of the classification error (sensitivity or TPR) as a function of the considered link density, for three different scenarios: the use of only the raw features, as described in Table 3 (solid black squares); the use of the features extracted from the parenclitic representation alone (hollow black circles); and the use of the combined sets of features (solid blue triangles). Note that, in the former case, the result is constant, as the original features are not affected by the binarization process. In order to avoid overfitting, this classification has been performed on a balanced sub data set, composed of an equal number of legal and illegal transactions.

Several conclusions can be drawn from Figure 2. First of all, the features extracted from the parenclitic networks are not enough, alone, to reach a low classification error. This has to be expected: while important information can be codified in the interaction between raw features, some
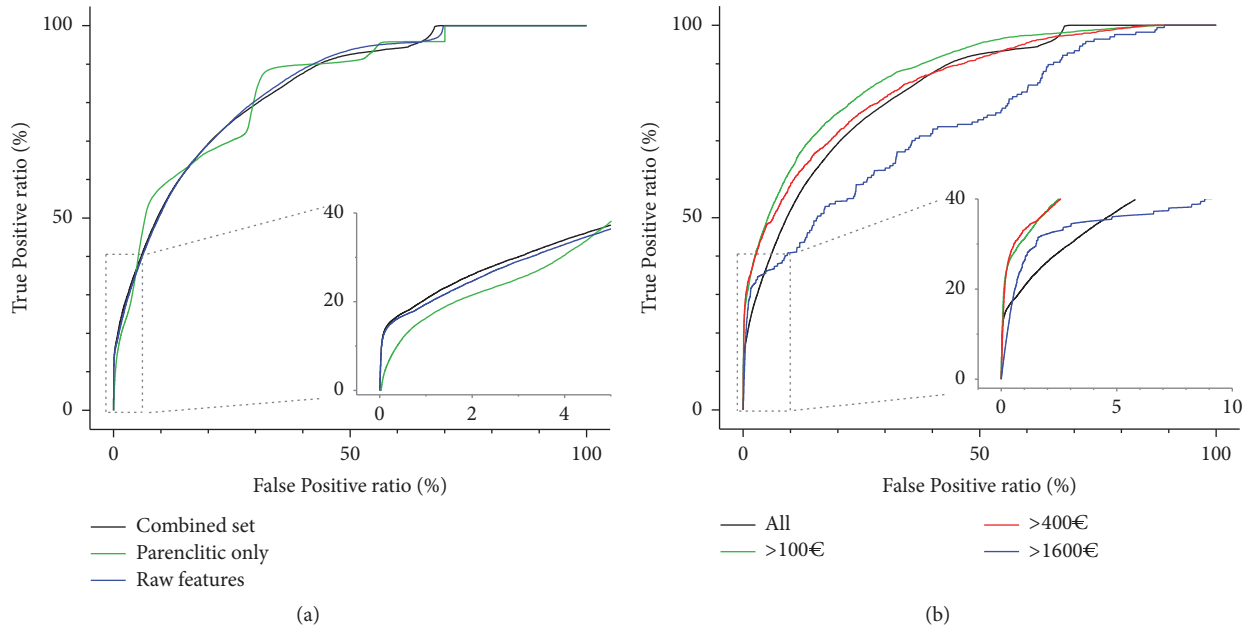
FIGURE 3: (a) ROC curves of the classification, corresponding to the use of the raw features alone (blue line), of the parenclitic features (green line), and of the combined sets (black line). (b) ROC curves, obtained through the combined data set, as a function of the transaction sizes.

important clues may be hidden in the latter, e.g., abnormal transaction sizes or timings. At the same time, the addition of parenclitic features to the raw data set enhances the obtained results, with the error dropping from a 19.2% to a 12.23%. This is further illustrated in Figure 2(b), depicting the reduction in the classification error (in percentage) when considering only parenclitic features and the whole data set; note that, in the first case, the reduction is negative as the error increases. Finally, the best classification suggests that the optimal link density that should be considered is of 60%, meaning that the 40% of links with less weight should be deleted.

If Figure 2 is useful to detect the best link density for the analysis, it does not convey information about the real performance of the classification algorithm in an operational environment. For that, Figure 3(a) presents three ROC curves, corresponding to the use of raw (blue line), parenclitic (green line), and combined features (black line) as before. Note that results here presented correspond to the optimal link density of 60%, as previously estimated. As previously discussed, the most interesting operational configuration is the one minimizing the number of false positives, as this minimizes the commercial costs of the organization. The inset of Figure 3 thus shows the bottom left part of the curves. It can be appreciated that, after an initial part in which results are comparable, the addition of the parenclitic features slightly increases the number of true positives; note how the black line is above the blue one.

In order to confirm such graphical results, two Areas Under the ROC Curve (AUC) have been calculated, respectively, corresponding to the whole graph and for ratios of False Positive below 4% (and thus for the part of the graph included in the inset). In the first case, results are 0.8365, 0.8392, and 0.8388, respectively, for the combined set, parenclitic only features, and raw features; while, in the second case, resulting AUCs are 0.2550, 0.2026, and 0.2407. If results are roughly comparable in the global behavior, an interesting increase in the AUC (of a 5.9%) is observed when one's objective is the minimization of the number of false positives.

Even though this may seem a negligible difference, it is worth noting that any improvement, however small, has a significant impact due to the large number of transactions managed by the system. Increasing the fraud detection rate by 1% would allow identifying ≈20.000 new illicit transactions per year, or ≈2 M€ in saved costs.

As well known in data mining, the result of the comparison of the two models of interest (here, the classification with the raw data alone, and with the combined data set) should be done with special care, as the number of features in both sets is different, respectively 10 and 17. This may lead to an overfitting of the latter model, and thus to the obtention of higher-than-real scores [53]. It must be noted that overfitting is not expected to be an issue here, firstly because the training and evaluation sets are kept separated, such that no information about the actual results is explicitly encoded in the training data; and secondly, because of the high number fo instances ($1.8 \cdot 10^8$) used in the evaluation. Additionally, 100 classifications have been performed on synthetic data sets, composed of the raw features and of a random permutation of the parenclitic networks' metrics. The resulting AUC (0.2387 ± 0.0083) corresponds to a Z-Score of +1.723, thus indicating that the obtained result is statistically significant for a significance level of 0.05.
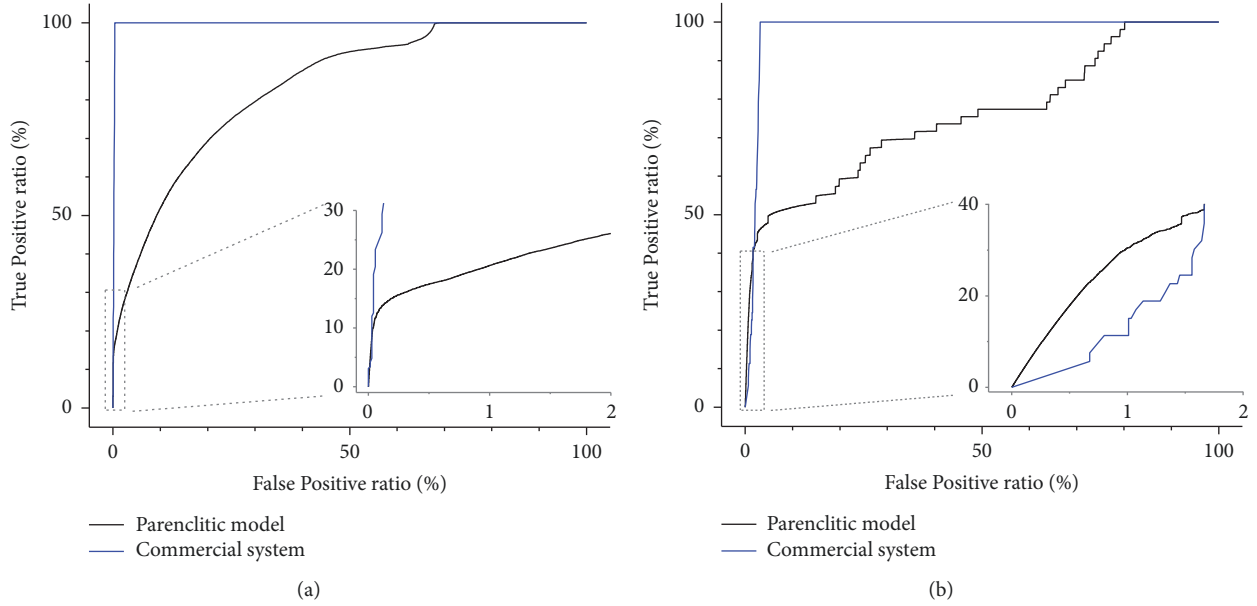
FIGURE 4: (a) ROC curves for the network-based model (black line) and a commercial system (blue line). (b) ROC curves, for the proposed network-based algorithm and a commercial system, when only on-line (Internet) transactions are considered.

Figure 3(b) further presents four ROC curves calculated for different transaction sizes: all transactions (black line), and transactions above 100€ (green line, AUC = 0.3835 for False Positive ratios between 0% and 5%), 400€ (red line, AUC = 0.4040), and 1.600€ (blue line, AUC = 0.2989). Deleting small transactions results in an improvement of the detection efficiency; note how the green and red lines lay above the black one. Additionally, the proposed algorithm fails for large transactions; this does not come as a surprise, as the larger the size, the fewer the available instances, making training more challenging.

If what previously presented illustrates that the use of a network representation can improve a fraud detection algorithm, it does not clarify how it ranks against a commercial system. As may be expected, the proposed algorithm is less efficient than the fraud score included in the original data set; see Figure 4(a) (due to confidentiality issues, the name and characteristics of the commercial fraud detection system cannot be included in this publication). Nevertheless, there are niches in which the opposite happens, the most important being the analysis of on-line transactions. Figure 4(b) depicts two ROC curves, respectively, for the algorithm based on parenclitic networks (black line) and the commercial system (blue line), when only transactions realized through Internet are considered. While the commercial system clearly outperforms the proposed algorithm, with an Area Under the Curve (AUC) close to 1.0, the latter is slightly better for a low ratio of False Positive, as previously explained, the plane region most interesting for real operations.

5. Conclusions

Complex networks and data mining models share more characteristics than what we could have expected in the first naive approach, most notably having similar objectives: both aim at extracting information from (potentially complex) systems to ultimately generate new compact quantifiable representations. At the same time, they approach this common problem from two different approaches: the former by extracting and quantitatively evaluating the underlying structure; the latter by creating predictive models based on historical data [11]. In this paper we test the hypothesis that complex networks can be used as a way to improve data mining models, framed within the problem of detecting fraud instances in credit card transactions, providing a new example about how complex networks and data mining may be integrated as complementary tools in a synergistic manner in order to improve the classification rates obtained by classical data mining algorithms.

Results confirm that features extracted from a network-based representation of data, leveraging on a recently proposed parenclitic approach [46, 47], can play an important role: while not effective in themselves, such features can improve the score obtained by a standard ANN classification model. We further show how the resulting model is especially efficient in detecting frauds in some niches of operations, like medium-sized and on-line transactions. Finally, we illustrate as in the latter case that the network-based model is able to yield better results than a commercial fraud detection system. All results have been obtained with a unique data set, comprising all transactions managed during two years by a major Spanish bank and including more than 180 million operations.

Conflicts of Interest

The authors declare no financial conflicts of interest.

Authors' Contributions

Acknowledgments

References

[1] E. W. T. Ngai, Y. Hu, Y. H. Wong, Y. Chen, and X. Sun, "The application of data mining techniques in financial fraud detection: a classification framework and an academic review of literature," *Decision Support Systems*, vol. 50, no. 3, pp. 559–569, 2011.

[2] J. West and M. Bhattacharya, "Intelligent financial fraud detection: a comprehensive review," *Computers & Security*, vol. 57, pp. 47–66, 2016.

[3] T. P. Bhatla, V. Prabhu, and A. Dua, "Understanding credit card frauds," *Cards Business Review*, vol. 1, 2003.

[4] J. Rollins and C. Wilson, *Terrorist capabilities for cyberattack: Overview and policy issues*, 2006.

[5] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning*, vol. 1, Springer, Berlin, Germany, 2001.

[6] J. Han, J. Pei, and M. Kamber, *Data mining: concepts and techniques*, Elsevier, 2011.

[7] I. H. Witten, E. Frank, M. A. Hall, and C. J. Pal, *Data Mining: Practical Machine Learning Tools and Techniques*, Morgan Kaufmann, 2016.

[8] S. H. Strogatz, "Exploring complex networks," *Nature*, vol. 410, no. 6825, pp. 268–276, 2001.

[9] R. Albert and A. Barabási, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, no. 1, pp. 47–97, 2002.

[10] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D. W. Hwang, "Complex networks: Structure and dynamics," *Physics Reports*, vol. 424, no. 4-5, pp. 175–308, 2006.

[11] M. Zanin, D. Papo, P. A. Sousa et al., "Combining complex networks and data mining: why and how," *Physics Reports*, vol. 635, pp. 1–44, 2016.

[12] S. Ghosh and D. L. Reilly, "A tutorial on hidden markov models and selected applications in speech recognition," in *Proceedings of the 27th Hawaii International Conference on System Sciences: Information Systems: Decision Support and Knowledge-Based Systems*, vol. 3, pp. 621–630, 1994.

[13] E. Aleskerov, B. Freisleben, and B. Rao, "Cardwatch: A neural network based database mining system for credit card fraud detection," in *Proceedings of the 1997 IEEE/IAFE Conference on Computational Intelligence for Financial Engineering, CIFEr*, pp. 220–226, IEEE, March 1997.

[14] R. Brause, T. Langsdorf, and M. Hepp, "Neural data mining for credit card fraud detection," in *Proceedings of the 11th IEEE International Conference on Tools with Artificial Intelligence*, pp. 103–106, IEEE, 1999.

[15] S. Maes, K. Tuyls, B. Vanschoenwinkel, and B. Manderick, "Credit card fraud detection using Bayesian and neural networks," in *Proceedings of the 1st International Naiso Congress on Neuro Fuzzy Technologies*, pp. 261–270, 2002.

[16] M. Syeda, Y.-Q. Zhang, and Y. Pan, "Parallel granular neural networks for fast credit card fraud detection," in *Proceedings of the IEEE International Conference on Fuzzy Systems (FUZZ-IEEE '02)*, vol. 1, pp. 572–577, 2002.

[17] R. Patidar and L. Sharma, "Credit card fraud detection using neural network," *International Journal of Soft Computing and Engineering (IJSCE)*, vol. 1, 2011.

[18] E. M. Carneiro, L. A. V. Dias, A. M. D. Cunha, and L. F. S. Mialaret, "Cluster analysis and artificial neural networks: a case study in credit card fraud detection," in *Proceedings of the 12th International Conference on Information Technology: New Generations, ITNG 2015*, pp. 122–126, IEEE, April 2015.

[19] C. Serrano-Cinca, "Self organizing neural networks for financial diagnosis," *Decision Support Systems*, vol. 17, no. 3, pp. 227–238, 1996.

[20] V. Zaslavsky and A. Strizhak, "Credit card fraud detection using self organizing maps," *The International Journal of Information & Security*, vol. 18, pp. 48–63, 2006.

[21] P. J. Bentley, J. Kim, G.-H. Jung, and J.-U. Choi, "Fuzzy darwinian detection of credit card fraud," in *Proceedings of the 14th Annual Fall Symposium of the Korean Information Processing Society*, vol. 14, pp. 1–4, 2000.

[22] E. Duman and M. H. Ozcelik, "Detecting credit card fraud by genetic algorithm and scatter search," *Expert Systems with Applications*, vol. 38, no. 10, pp. 13057–13063, 2011.

[23] K. RamaKalyani and D. UmaDevi, "Fraud detection of credit card payment system by genetic algorithm," *International Journal of Scientific & Engineering Research*, vol. 3, no. 7, pp. 1–6, 2012.

[24] R. Chen, T. Chen, Y. Chien, and Y. Yang, *Advances in Neural Networks–ISNN 2005 821*, 2005.

[25] R.-C. Chen, T.-S. Chen, and C.-C. Lin, "A new binary support vector system for increasing detection rate of credit card fraud," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 20, no. 2, pp. 227–239, 2006.

[26] Y. G. Sahin and E. Duman, "Detecting credit card fraud by decision trees and support vector machines," in *Proceedings of the International Multi-Conference of Engineering and Computer Statistics*, vol. 1, 2011.

[27] Q. Lu and C. Ju, "Research on credit card fraud detection model based on class weighted support vector machine," *Journal of Convergence Information Technology*, vol. 6, no. 1, pp. 62–68, 2011.

[28] L. D. F. Costa, F. A. Rodrigues, G. Travieso, and P. R. V. Boas, "Characterization of complex networks: a survey of measurements," *Advances in Physics*, vol. 56, no. 1, pp. 167–242, 2007.

[29] B. Wang, H. Tang, C. Guo, and Z. Xiu, "Entropy optimization of scale-free networks' robustness to random failures," *Physica A: Statistical Mechanics and its Applications*, vol. 363, no. 2, pp. 591–596, 2006.

[30] V. Latora and M. Marchiori, "Efficient behavior of small-world networks," *Physical Review Letters*, vol. 87, Article ID 198701, 2001.

[31] M. Zanin, P. A. Sousa, and E. Menasalvas, "Information content: Assessing meso-scale structures in complex networks," *Europhysics Letters*, vol. 106, article 30001, 2014.

[32] S. Bhattacharyya, S. Jha, K. Tharakunnel, and J. C. Westland, "Data mining for credit card fraud: a comparative study," *Decision Support Systems*, vol. 50, no. 3, pp. 602–613, 2011.

[33] N. Sethi and A. Gera, "A revived survey of various credit card fraud detection techniques," *International Journal of Computer Science and Mobile Computing*, vol. 3, no. 4, pp. 780–791, 2014.

[34] Z. Zojaji, R. E. Atani, and A. H. Monadjemi, "A survey of credit card fraud detection techniques: data and technique oriented perspective," *Cryptography and Security*, 2016.

[35] J. M. Zurada, *Introduction to Artificial Neural Systems*, West St. Paul, 1992.

[36] M. T. Hagan, H. B. Demuth, and M. H. Beale, *Neural Network Design*, Pws Pub., Boston, Mass, USA, 1996.

[37] L. Davis, Ed., *Handbook of Genetic Algorithms*, Van Nostrand Reinhold, New York, NY, USA, 1991.

[38] M. Kumar, M. Husian, N. Upreti, and D. Gupta, "Genetic algorithm: Review and application," *International Journal of Information Technology and Knowledge Management*, vol. 2, pp. 451–454, 2010.

[39] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[40] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines And Other Kernel-Based Learning Methods*, Cambridge university press, 2000.

[41] F. V. Jensen, *An Introduction to Bayesian Networks*, vol. 210, UCL press, London, UK, 1996.

[42] N. Friedman, D. Geiger, and M. Goldszmidt, "Bayesian network classifiers," *Machine Learning*, vol. 29, no. 2-3, pp. 131–163, 1997.

[43] T. Kohonen, "The self-organizing map," *Neurocomputing*, vol. 21, no. 1–3, pp. 1–6, 1998.

[44] M. F. A. Gadi, X. Wang, and A. P. do Lago, "Credit card fraud detection with artificial immune system," in *ICARIS 2008: Artificial Immune Systems*, vol. 8 of *Lecture Notes in Computer Science*, pp. 119–131, Springer, 2008.

[45] J. A. Hanley and B. J. McNeil, "The meaning and use of the area under a receiver operating characteristic (ROC) curve," *Radiology*, vol. 143, no. 1, pp. 29–36, 1982.

[46] M. Zanin and S. Boccaletti, "Complex networks analysis of obstructive nephropathy data," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 21, no. 3, 033103, 5 pages, 2011.

[47] M. Zanin, J. M. Alcazar, J. V. Carbajosa et al., "Parenclitic networks: Uncovering new functions in biological data," *Scientific Reports*, vol. 4, article no. 5112, 2014.

[48] M. Zanin, P. Sousa, D. Papo et al., "Optimizing functional network representation of multivariate time series," *Scientific Reports*, vol. 2, article 630, 2012.

[49] M. Zanin, D. Papo, M. Romance, R. Criado, and S. Moral, "The topology of card transaction money flows," *Physica A: Statistical Mechanics and its Applications*, vol. 462, pp. 134–140, 2016.

[50] F. Rosenblatt, "The perceptron: a probabilistic model for information storage and organization in the brain," *Psychological Review*, vol. 65, no. 6, pp. 386–408, 1958.

[51] P. J. Werbos, *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*, secondoftwo school Harvard University, 1974.

[52] M. R. Berthold, N. Cebron, F. Dill et al., "KNIME - the Konstanz information miner," *ACM SIGKDD Explorations Newsletter*, vol. 11, no. 1, pp. 26–31, 2009.

[53] D. M. Hawkins, "The problem of overfitting," *Journal of Chemical Information and Computer Sciences*, vol. 44, no. 1, pp. 1–12, 2004.