

Available online at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/coseComputers
&
Security

Intelligent financial fraud detection: A comprehensive review

Jarrod West, Maumita Bhattacharya *

School of Computing & Mathematics, Charles Sturt University, Albury, NSW 2640, Australia

ARTICLE INFO

Article history:

Received 11 September 2014

Received in revised form 10 April 2015

Accepted 8 September 2015

Available online 12 November 2015

Keywords:

Financial fraud detection

Computational intelligence

Data mining

Anomaly detection

Classification

ABSTRACT

Financial fraud is an issue with far reaching consequences in the finance industry, government, corporate sectors, and for ordinary consumers. Increasing dependence on new technologies such as cloud and mobile computing in recent years has compounded the problem. Traditional methods involving manual detection are not only time consuming, expensive and inaccurate, but in the age of big data they are also impractical. Not surprisingly, financial institutions have turned to automated processes using statistical and computational methods. This paper presents a comprehensive review of financial fraud detection research using such data mining methods, with a particular focus on computational intelligence (CI)-based techniques. Over fifty scientific literature, primarily spanning the period 2004–2014, were analysed in this study; literature that reported empirical studies focusing specifically on CI-based financial fraud detection were considered in particular. Research gap was identified as none of the existing review articles addresses the association among fraud types, CI-based detection algorithms and their performance, as reported in the literature. We have presented a comprehensive classification as well as analysis of existing fraud detection literature based on key aspects such as detection algorithm used, fraud type investigated, and performance of the detection methods for specific financial fraud types. Some of the key issues and challenges associated with the current practices and potential future direction of research have also been identified.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Financial fraud is an issue that has wide reaching consequences in both the finance industry and daily life. Fraud can reduce confidence in industry, destabilise economies, and affect people's cost of living. Traditional approaches relied on manual techniques such as auditing, which are inefficient and unreliable due to the difficulty of the problem. Data mining-based approaches have been shown to be useful because of their ability to identify small anomalies in large data sets (Ngai et al., 2011). There are many different types of fraud, as well as a

variety of data mining methods, and research is continually being undertaken to find the best approach for each case.

Financial fraud is a broad term with various potential meanings, but for our purposes it can be defined as the intentional use of illegal methods or practices for the purpose of obtaining financial gain (Zhou and Kapoor, 2011). Fraud has a large negative impact on business and society: credit card fraud alone accounts for billions of dollars of lost revenue each year (Bhattacharyya et al., 2011), and some figures suggest that the total yearly cost to the US could be in excess of \$400 billion (Kirkos et al., 2007), while a third study shows that UK insurers are out 1.6 billion pounds a year due to fraudulent claims

* Corresponding author. Tel.: +61 2 60519619.

E-mail address: maumita.bhattacharya@ieee.org (M. Bhattacharya).
<http://dx.doi.org/10.1016/j.cose.2015.09.005>

0167-4048/© 2015 Elsevier Ltd. All rights reserved.

(Ngai et al., 2011). Financial fraud also has broader ramifications on the industry, such as providing funding for illicit activities like drug trafficking and organised crime (Bhattacharyya et al., 2011). For credit card fraud the cost is typically worn by the merchants, who end up paying shipping, chargeback, and administrative costs as well as losing consumer confidence after being victim to a fraudulent transaction (Quah and Sriganesh, 2008; Sánchez et al., 2009). In this way we can see the widespread consequences that fraud can have and the importance in minimising it.

Advancements in modern technologies such as the internet and mobile computing have led to an increase in financial fraud in recent years (Yeh and Lien, 2009). Social factors such as the increased distribution of credit cards has not only increased spending but also resulted in an increase to fraud (Sánchez et al., 2009). Fraudsters are continually refining their methods, and as such there is a requirement for detection methods to be able to evolve accordingly (Bhattacharyya et al., 2011). Data mining has already been shown to be useful in similar domains such as credit card approval, bankruptcy prediction, and analysis of share markets (Panigrahi et al., 2009). Fraud detection is considered to be a similar classification problem but with a vast imbalance in fraudulent to legitimate transactions, and a sizeable difference in cost for misclassifying them (Duman and Ozcelik, 2011). Data mining approaches are also applicable to fraud detection for their efficiency at processing large datasets and their ability to work without requiring knowledge of the input variables (Ravisankar et al., 2011).

A useful framework for applying data mining to fraud detection is to use it as a method for classifying suspicious transactions or samples for further consideration. Studies show that reviewing 2% of credit card transactions could reduce fraud losses to 1% of the total cost of all purchases, with more assessments resulting in smaller loss but with an increase in auditing costs (Quah and Sriganesh, 2008). A multi-layer pipeline approach can be used with each step applying a more rigorous method to detect fraud. Data mining can be utilised to efficiently filter out more obvious fraud cases in the initial levels and leave the more subtle ones to be reviewed manually (Quah and Sriganesh, 2008).

In this article we will use a few broad terminologies that are defined here for clarity. Data mining refers to any method that processes large quantities of data to derive an underlying meaning. Within this classification we will consider two categories of data mining: statistical and computational. We define the statistical techniques as those that are based on traditional mathematical methods, such as logistic regression and Bayesian theory. Computational methods are those which use modern intelligence techniques, such as neural networks and support vector machines. Though these categories share many similarities, we will consider that the main difference between them is that computational methods are capable of learning from and adapting to the problem domain, while statistical methods are more rigid. Both types of data mining will be researched in this article.

The objective of this paper is to provide a review of existing literature in financial fraud detection and compare their findings (focussing primarily on literature published during 2004–2014 and research that reports empirical studies using

CI-based techniques). A key focus of this review is on the reported performance of CI techniques for specific fraud types. None of the existing reviews (for example Ngai et al., 2011, Phua et al., 2010, and Edge and Sampaio, 2009) covers this aspect. This will provide an indication to future researchers on the areas that are currently available for future study. The remainder of the article is structured as follows. In the next section we will detail the history of intelligent fraud detection research. Section 3 will list the different types of financial fraud. Section 4 presents an overview of the assorted detection methods used. Section 5 details the specific efforts of previous researchers in detecting financial fraud. Section 6 offers an insight into what is missing from existing techniques and proposes areas for future research. Section 7 provides a conclusion to our research and a discussion of our findings.

2. Related work

Prior investigation has already been performed on some aspects of intelligent financial fraud detection, and a brief history of the specific research methods undertaken is given here. Fig. 1 shows the timeline of research for the last decade.

Initial fraud detection studies focussed heavily on statistical models such as logistic regression, as well as neural networks (Quah and Sriganesh, 2008; Yue et al., 2007; Hoogs et al., 2007). Zhang et al. discovered that neural networks had been used for financial applications such as forecasting since 1988 (Zhang et al., 1998). In 1995 Sohl and Venkatchalam first predicted financial statement fraud using a back-propagation neural network (Sohl and Venkatchalam, 1995). Fraser et al. compared techniques across a quantitative spectrum including statistical and computational methods such as regression and neural networks (Fraser et al., 1997). In 1998 Fanning and Cogger used a neural network based on other financial ratios and variables and found it compared favourably to discriminant analysis and logistic regression (Fanning and Cogger, 1998). In 2001 and 2002 Bolton and Hand performed some general analysis on fraud detection, focussing specifically on statistical learning (Bolton and Hand, 2002; Bolton and Hand, 2001), and Rezaee investigated financial statement fraud in depth (Rezaee, 2002). Table 1 shows the bias towards neural networks and statistical methods with a list of further fraud detection research performed outside of the last decade.

Recent fraud detection research has been far more varied in methods studied, although the former techniques are still popular. In 2004 Kou et al. reviewed the study of general fraud detection using analytic techniques including neural networks (Kou et al., 2004). Vatsa et al. investigated a novel approach using game theory in 2005, which modelled fraudsters and detection methods as opposing players in a game, each striving to obtain the greatest financial advantage (Vatsa et al., 2005). The following year Yang and Hwang studied health care fraud using a process-mining approach (Yang and Hwang, 2006).

In 2007 Pinquet et al. and Viaene et al. both studied logistic regression with insurance fraud, concentrating on a database of Spanish automobile insurance claims (Pinquet et al., 2007; Viaene et al., 2007). Kirkos et al. compared statistical methods with neural networks to identify fraudulent Greek manufacturing companies (Kirkos et al., 2007), and Bose and Wang

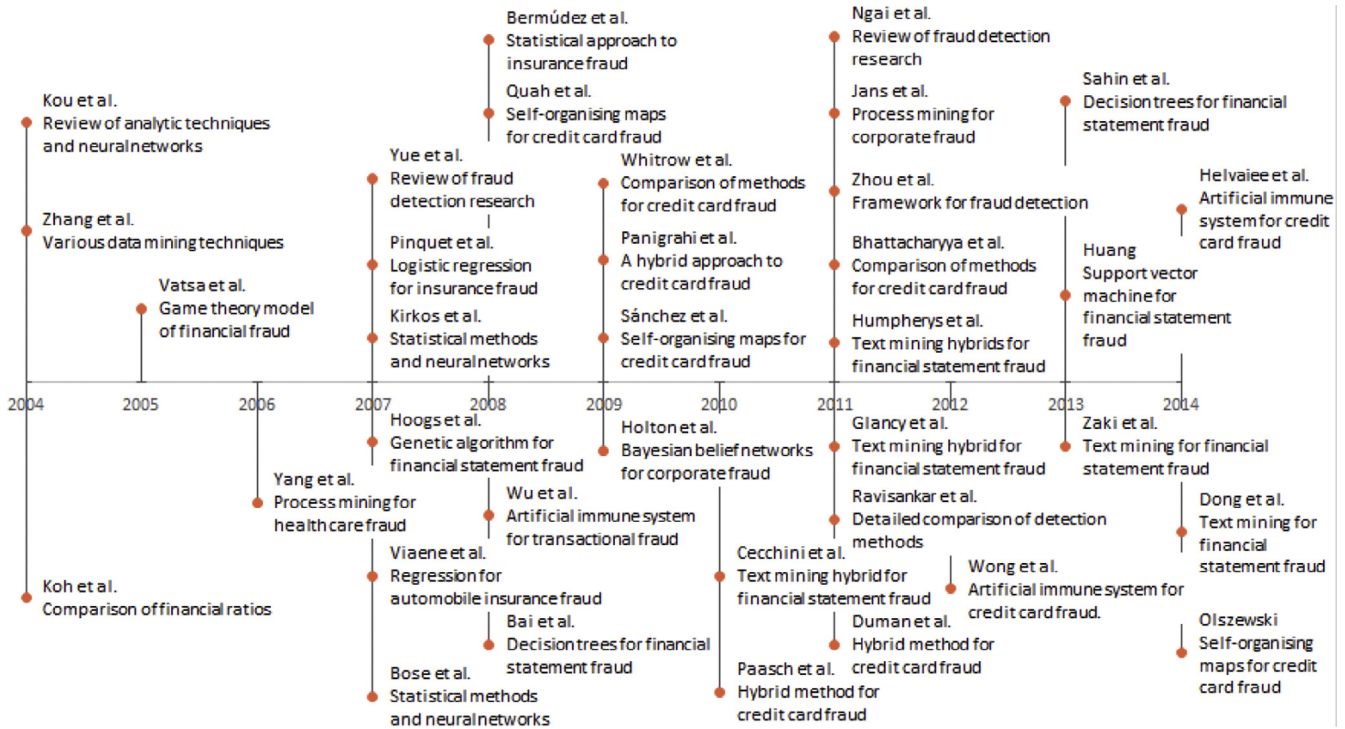


Fig. 1 – Recent financial fraud detection research.

focused on classification and regression trees to solve financial statement fraud in a selection of Chinese companies (Bose and Wang, 2007).

Also in 2007 Hoogs et al. introduced a genetic algorithm on Accounting and Auditing Enforcement Releases to detect fraudulent companies in the US (Hoogs et al., 2007), and Yue et al. performed a review of existing fraud detection literature. They claimed that the only successful methods of fraud detection to date, as well as the most commonly researched, were classification based (Yue et al., 2007).

Bai et al. used decision trees to study financial statement fraud for a selection of Chinese companies in 2008 (Bai et al., 2008). Bermúdez et al. took a statistical approach to insurance fraud detection, using the same samples that Pinquet et al. and Viaene et al. used previously (Bermúdez et al., 2008). Quah and Sriganesh looked at visualising credit card fraud with self-organising maps, focussing on real-world samples from the Singaporean branch of an international bank (Quah and Sriganesh, 2008). Wu and

Banzhaf modified the standard artificial immune system method with a coevolutionary approach, using it to solve transactional fraud with automatic teller and point-of-sale data for a financial institution (Wu and Banzhaf, 2008).

In 2009 Holton utilised a combination of text mining and Bayesian belief networks to identify disgruntled employees likely to commit corporate fraud (Holton, 2009). Panigrahi et al. combined a Dempster–Schaefer adder with a Bayesian learner to solve credit card fraud with their own synthesised data (Panigrahi et al., 2009). Sánchez et al. focussed on credit cards provided by a multinational department store, using self-organising maps to cluster and visualise fraudulent patterns (Sánchez et al., 2009). Whitrow et al. compared support vector machines with decision trees in solving credit card fraud, with a focus on aggregating common transactional variables to create new inputs (Whitrow et al., 2009).

In 2010 Cecchini et al. studied Accounting and Auditing Enforcement Releases (AAER) with their own text mining and

Table 1 – Examples of research on fraud detection prior to 2004.

Method investigated	Year of publication	Research
Neural network	1997	Aleskerov et al. (1997); Green and Choi (1997); Busta and Weinberg (1998); Feroz et al.
	1998	(2000); Calderon and Cheh (2002)
	2000	
	2002	
	2002	
Regression	1998	Summers and Sweeney, (1998); Beneish (1999); Abbott et al., 2000; Bell and Carcello
	1999	(2000); Spathis et al. (2002); Spathis (2002)
	2000	
	2002	
	2002	
Fuzzy logic	1998	Deshmukh and Talluru (1998)
Other statistical methods	1996	Hansen et al. (1996); Stolfo et al. (1997)
	1997	

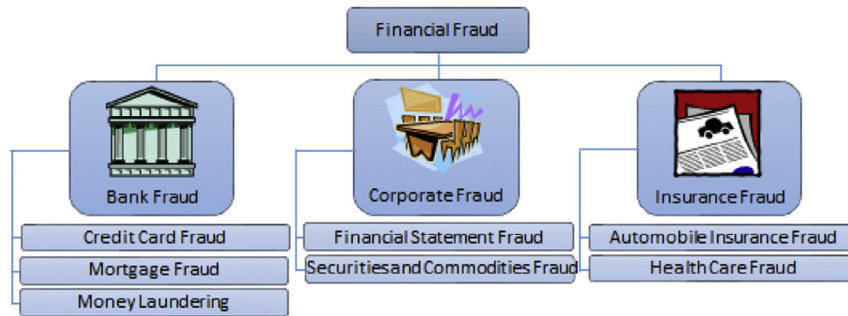


Fig. 2 – Common types of financial fraud.

support vector machine hybrid to predict financial statement fraud in US companies (Cecchini et al., 2010). Then in 2011 Bhattacharyya et al. compared the ability of logistic regression, support vector machines, and random forests on a large sample of credit card transactions to identify which were fraudulent (Bhattacharyya et al., 2011). Duman and Ozcelik combined the strengths of genetic algorithms and scatter search to create their own hybrid method. They used it to track consumer spending with a large Turkish bank, as an aid to predicting the occurrence of credit card fraud (Duman and Ozcelik, 2011).

Humphrys et al. created text mining hybrids by utilising other common methods to act as the classifier. With a support vector machine, decision tree, and Bayesian belief network they managed to successfully identify fraud within the company's 10-K document filings (Humpherys et al., 2011). Glancy and Yadav also studied sections of 10-K documents for US companies known to be fraudulent, processing the text with a singular validation decomposition vector to classify the samples (Glancy and Yadav, 2011). Jans et al. applied process mining to the internal logs created by a European financial institution to detect corporate fraud (Jans et al., 2011), and Ngai et al. performed a substantial review of existing fraud detection literature as of 2011 (Ngai et al., 2011).

Also in 2011 Ravisankar et al. compared a large range of methods to identify financial statement fraud within Chinese companies. In addition to support vector machines they looked at genetic programming, logistic regression, group method of data handling, and variety of neural networks (Ravisankar et al., 2011). Zhou and Kapoor created a generic framework for financial statement fraud detection using response surface methodology (Zhou and Kapoor, 2011), then in 2012 Wong et al. utilised an artificial immune system to predict credit card fraud for a major Australian bank (Wong et al., 2012).

In 2013 Huang investigated financial statement fraud in a series of Taiwanese companies using logistic regression and a support vector machine (Huang, 2013). Zaki and Theodoulidis took a more direct approach and focussed on the litigation section of the Securities and Exchange Commission website, applying their own text mining algorithm to classify financial statement fraud (Zaki and Theodoulidis, 2013). Sahin et al. studied the ability of decision trees to identify fraudulent credit card transactions, using a 6-month sample from a major bank (Sahin et al., 2013).

In 2014 Dong et al. used text mining to study AAERs for Chinese companies that were trading publicly in the US (Dong

et al., 2014). Olszewski visualised credit card fraud with self organising maps, focussing only on accounts held by residents of Warsaw, Poland (Olszewski, 2014); Soltani Halvaiee and Akbari utilised an artificial immune system to identify credit card fraud for an anonymous Brazilian bank (Soltani Halvaiee and Akbari, 2014); and West et al. investigated the present state of fraud detection research (West et al., 2014).

3. Types of financial fraud

There are many different types of financial fraud, and a brief description of some of the major types will be listed here (also see Fig. 2). The fraud types were selected from the list provided in the Federal Bureau of Investigation, Financial Crimes Report (2010–2011), United States (FBI, 2010–2011).

3.1. Credit card fraud

Credit card fraud refers to the unauthorised use of a person's credit card to perform fraudulent transactions without the user's knowledge (Ngai et al., 2011). The transactions can be performed using the physical card, where the card was either lost or stolen, but is often performed remotely (Bhattacharyya et al., 2011). The cardholder's information may be acquired by one of a few methods. Phishing involves a fraudster impersonating a finance official to convince the user to divulge their details, swipers or skimmers provide an interface to an ATM or POS device which can read the card directly, or entire databases of user's information can be obtained if the fraudster is able to breach the financial institutions network security or enlist the help of an accomplice within the company (Bhattacharyya et al., 2011; Quah and Sriganesh, 2008). Obtaining the user's card could even be as simple as intercepting mail containing a new or replacement card (Quah and Sriganesh, 2008). The anonymity and availability of these remote methods have given rise to the prevalence of organised crime in credit card fraud (Bhattacharyya et al., 2011). A typical method for identifying credit card fraud is to analyse a customer's regular spending habits and flag transactions which are noticeably outside of this model (Duman and Ozcelik, 2011).

3.2. Securities and commodities fraud

Securities fraud, also known as commodities fraud, refers to a variety of methods by which a person is deceived into in-

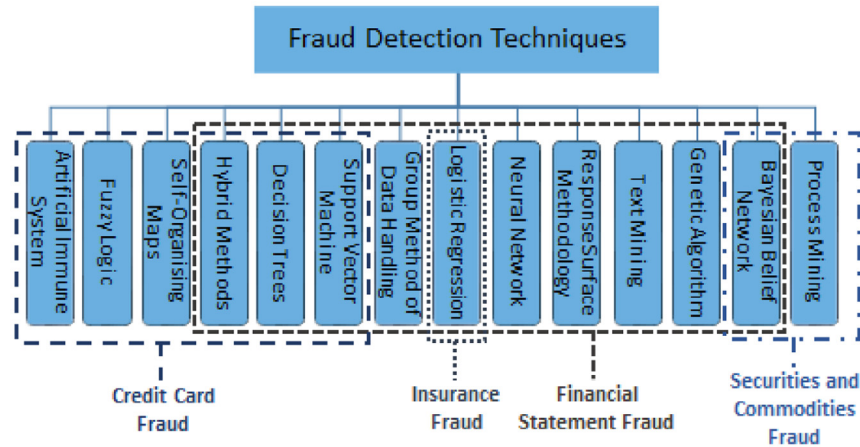


Fig. 3 – Detection methods and fraud types.

vesting into a company based on false information. It includes Pyramid Schemes, Ponzi Schemes, Hedge Fund Fraud, Foreign Exchange Fraud, and Embezzlement (Ngai et al., 2011).

3.3. Financial statement fraud

Financial statements are the documents release by a company that explain details such as their expenses, loans, income and profits (Ravisankar et al., 2011). They can also include comments from management on the business' performance and expected issues that may arise in the future (Glancy and Yadav, 2011). The various financial statements that the company releases give an overall picture of the company's status, and can be used to indicate how successful the company is, influence stock prices, and determine if they are applicable for loans (Ravisankar et al., 2011). Financial statement fraud, also known as corporate fraud, involves doctoring these statements to make the company appear more profitable.

Reasons for committing financial statement fraud include improving stock performance, reducing tax obligations, or as an attempt to exaggerate performance due to managerial pressure (Ravisankar et al., 2011). Financial statement fraud can be difficult to diagnose because of a general lack of understanding of the field, the infrequency in which it occurs, and the fact that it is usually committed by knowledgeable people within the industry who are capable of masking their deceit (Maes et al., 2002).

3.4. Insurance fraud

Insurance fraud is fraud that can be committed at any point during the insurance process, and by any people in the chain. Insurance claims fraud occurs when a customer submits a fraudulent insurance claim as a result of an exaggerated injury or loss of assets, or a completely fraudulent event. A common form of claims fraud is automobile insurance fraud, which is often committed by faking or intentionally committing accidents that result in excessive repair and injury costs. Larger scale claims fraud also occurs, such as crop insurance fraud where a consumer overstates their losses due to declining agricultural prices or the effects of natural disasters. Insurance

fraud can also include excessive billing, duplicate claims, kick-backs to brokers and "upcoding" of items (Ngai et al., 2011).

3.5. Mortgage fraud

Mortgage fraud is a specific form of financial fraud that refers to manipulation of a property or mortgage documents. It is often committed to misrepresent the value of a property for the purpose of influencing a lender to fund a loan for it (Ngai et al., 2011).

3.6. Money laundering

Money laundering is a method used by criminals to insert proceeds obtained from illicit ventures into valid businesses. This conceals the origin of the money, giving them the appearance of legitimate income and making it difficult to track their crimes. Money laundering is also undesirable as it enables the criminals to have economic influence (Ngai et al., 2011).

4. Computational intelligence and data mining methods for financial fraud detection

A large number of statistical and computational techniques exist that have been applied to data mining problems in recent years (see Fig. 3). This section contains a description of the operation of each method used in the reviewed literature, and Table 2 lists the relative strengths and limitations of each.

4.1. Bayesian belief networks

Bayesian belief networks are a statistical classification technique that makes use of Bayes theorem, a method to determine the probability that a given hypothesis is true. The theorem states that for a hypothesis H (such as whether an object X can be classified within a given class), the probability P is given by:

$$P(H|X) = \frac{P(X|H) * P(H)}{P(X)} \quad (1)$$

Table 2 – Relevant properties of data mining methods for financial fraud detection.

Method	Strengths	Limitations
Neural network	Well established history with fraud detection. Proven suitability with other non-algorithmic, binary classification problems.	Requires high computational power for training and operation, making it unsuitable for real-time function. Potential for overfitting if training set is not a good representation of the problem domain, so requires constant retraining to adapt to new methods of fraud.
Logistic model	Simple to implement. Well established history with fraud detection.	Lower classification performance than other data mining methods, difficulty with the complexity of fraud detection
Support vector machine	Capable of solving non-linear classification problems like fraud detection. Training and operation requires low computational power, which gives potential for real-time operation.	Difficult for auditors to process results due to transformation of input set
Decision trees, forests and CART	Simple to implement and understand. Training and operation requires low computational power, which gives potential for real-time operation.	Potential for overfitting if training set is not a good representation of the problem domain, so requires constant retraining to adapt to new methods of fraud. Optimisation during initial setup requires high computational power.
Genetic algorithm/programming	Simple to implement using classification accuracy as the fitness solution. Proven suitability with other non-algorithmic, binary classification problems.	Requires high computational power for training and operation, making it unsuitable for real-time function. Difficulty adapting to new fraud methods due to local maxima/minima problem.
Text mining	Highly useful for fraud types with large amounts of textual data, such as financial statement fraud.	Requires another classification method to perform the actual fraud detection. Textual data are typically more subjective, and thus harder to process.
Group method of data handling	Simple to implement. Guaranteed to provide the best available solution.	Difficulty classifying noisy data, which many fraud types contain.
Response-surface methodology	Capable of solving non-linear classification problems like fraud detection.	Lower classification performance than other data mining methods, difficulty with complexity of fraud detection.
Self-organising map	Simple to implement and very easy for auditors to understand given visual nature of results.	Visualisation requires auditor observation, cannot be fully automated easily.
Bayesian belief network	Proven suitability with other non-algorithmic, binary classification problems. High computational efficiency gives potential for real-time operation.	Requires strong understanding of typical and abnormal behaviour for the investigated fraud type.
Process mining	Useful for internal fraud investigations where information is available for every iterative step. Ability to focus on an entire process chain instead of individual attributes.	Requires strong understanding of typical and abnormal behaviour for the investigated fraud type. Difficulty classifying noisy data, which many fraud types contain.
Artificial immune system	High suitability for classification problems with imbalanced data, such as fraud detection.	Requires high computational power for operation, making it unsuitable for real-time function.
Hybrid methods	Adaptive to new fraud techniques by combining the strengths of multiple traditional detection methods.	Fraud is a high-cost problem and therefore new, under-tested methods carry a large amount of risk.

A Bayesian belief network uses a classifier to calculate $P(C_i|X)$ for all possible classes C_i and inserts X into the class with the highest $P(C_i|X)$. In this way the network is shown to classify each sample into the class that it is most likely to belong to (Kirkos et al., 2007).

Graphically, a Bayesian belief network can be modelled as a directed acyclic graph, with nodes to represent the samples and edges to reflect a causal dependency between them (see Fig. 4). Missing edges can then demonstrate that two variables are independent of one another (Ngai et al., 2011).

4.2. Logistic regression or logistic model

Logistic regression is a statistical method of classifying binary data which uses a linear model, often referred to as a logistic or logit model, to perform regression on a set of variables (Ngai

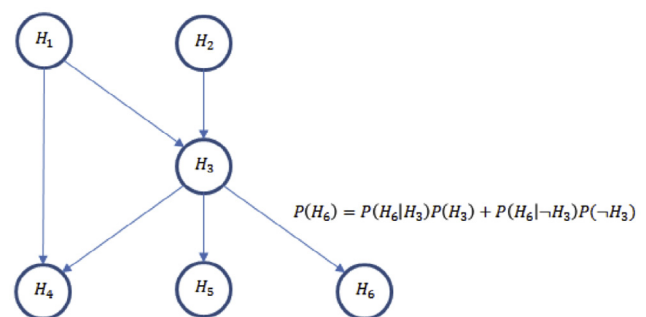


Fig. 4 – Example graphical representation of a Bayesian belief network showing the causal relationship between hypotheses H_6 and H_3 .

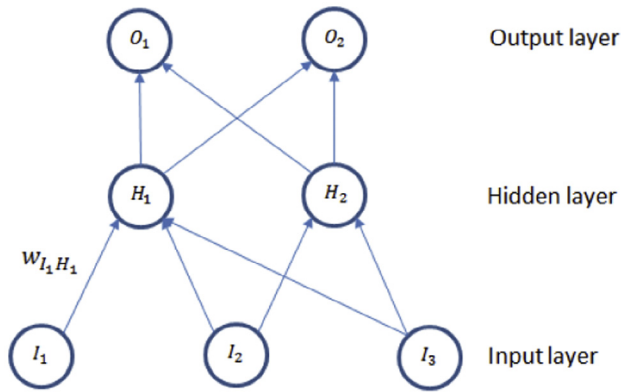


Fig. 5 – A simple neural network with three input variables, two outputs, and a single hidden layer.

et al., 2011; Ravisankar et al., 2011). It is a commonly used method for predicting patterns in data with unambiguous or numeric attributes (Ngai et al., 2011; Bhattacharyya et al., 2011). Logistic regression makes use of a series of input vectors and a dependent response variable, using the natural logarithm to calculate the probability that the result lies within a particular category. For binary classification the response variable takes the form:

$$y_i = \begin{cases} 1 \\ 0 \end{cases} \quad (2)$$

And the formula for calculating that a sample x_i belongs in class 1 is given by:

$$P(Y_i = 1 | X_i) = \frac{\exp(w_0 + w^T x_i)}{1 + \exp(w_0 + w^T x_i)} \quad (3)$$

where w_0 and w are the regression tuning parameters representing the intercept and coefficient vector respectively (Ravisankar et al., 2011).

4.3. Neural network

A neural network is a computational approximation of the human brain which uses a graph of vertices and edges to represent neurons and synapses (Ngai et al., 2011). The network performs by modelling the input variables as a layer of vertices and then assigning a weighting to each connection in the graph, while the other vertices are put into separate levels reflecting their distance from the input nodes (Kirkos et al., 2007) (see Fig. 5).

Each node considers its input as a function of the vertices connected to it at the previous layer. For each neuron j the signal received is given by:

$$u_j = \sum w_{ij} * x_i \quad (4)$$

where w_{ij} is the weight of the link between neurons i and j and x_i is the input of i . If the result is greater than a

predetermined amount the current neuron “fires” and becomes an input for the next layer (Kirkos et al., 2007).

A back propagation neural network is trained by running samples from a set of training data through the network and comparing the results. For the first iteration the weights at each edge are normally determined randomly, and after the results have been calculated each weighting is adjusted slightly for the next sequence (Zhang and Zhou, 2004). This continues until either the network has reduced its error to an acceptable amount or a predetermined iteration limit has been reached. After training the network’s performance may be tested with a set of validation data (Ngai et al., 2011). A common problem with back propagation neural networks is overtraining, which can cause the network to focus on tendencies particular to the training set but not the overall problem (Zhang and Zhou, 2004).

4.4. Support vector machine

Support vector machines are a classification method which works by converting a linear issue into a higher dimensional feature space. This enables complicated, non-linear problems such as financial fraud detection to be solved by linear classification without increasing the computational complexity. The function used to transform the dataset is called the kernel function, which can be considered as a mapping of points between the input space and a higher dimensional space. The kernel function is defined by:

$$k(x_1, x_2) = \langle \Phi(x_1), \Phi(x_2) \rangle \quad (5)$$

where $\Phi: X \rightarrow H$ maps points in input space X to higher-dimensional space H . After applying the kernel function to the dataset a hyperplane is used to separate the classes, which takes the form:

$$\langle w, x \rangle + b = 0 \quad (6)$$

The hyperplane is constructed in such a way as to maximise the separation between both classes, which helps to reduce potential errors caused by overtraining (see Fig. 6).

The classification for a support vector machine can therefore be defined as:

$$\sum_i \alpha_i y_i k(x_i, x) + b = 0 \quad (7)$$

The choice of which kernel function to use is dependent on the dataset and classification requirements, though there are many commonly used kernel functions such as the Gaussian radial basis function and polynomial function (Bhattacharyya et al., 2011).

4.5. Genetic algorithms and programming

Genetic algorithms use the concept of population evolution to iteratively improve solutions to the problem. It works by randomly creating a starting generation, then continuously reproducing each population using various techniques and choosing the survivors based on their fitness. Reproduction occurs by taking pairs of parents from the current generation

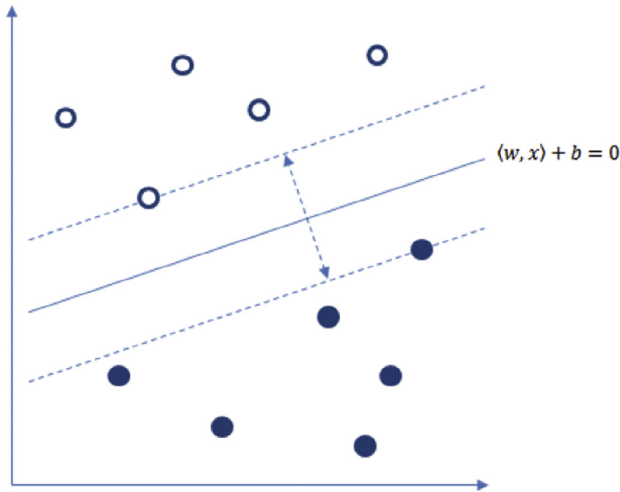


Fig. 6 – Two dimensional example showing the margin of separation between support vectors and hyperplane.

and applying crossover on two points, then randomly mutating a single element of the resulting child. The ability of the children is measured using a fitness function, and the result of this determines which parents and children are chosen to represent the next generation.

Measuring the fitness of the children can be as simple as measuring the percentage of samples they classify correctly. The algorithm terminates once a required fitness has been reached, but to avoid infinite looping a limit may be set on the number of iterations that are run (see Fig. 7). Genetic algorithms are similar to neural networks in that they require no prior knowledge of the problem domain and are capable of detecting underlying relationships between the samples (Ravisankar et al., 2011).

4.6. Decision trees, forests, and CART

Decision trees are a technique that classifies or predicts data using a tree with internal nodes representing binary choices on attributes and branches representing the outcome of that choice (Zhang and Zhou, 2004) (see Fig. 8). The nodes are created such that as a sample traverses the tree it is partitioned into subsets until it is eventually sorted into a mutually exclusive subgroup. Decision trees are also known as Classification and Regression Trees (CART) (Kirkos et al., 2007).

A decision forest, or random forest, is a collection of decision trees used to avoid the instability and risk of overtraining

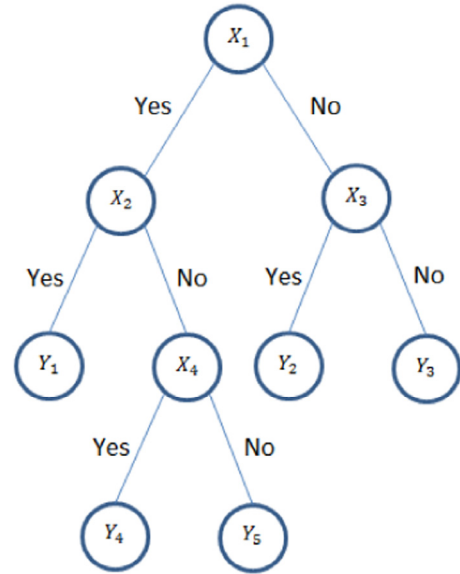


Fig. 8 – Binary decision tree where the input variables X_i will be a list of input predicates and the outputs Y_i will classify the results into one of multiple classes.

that can occur with a single tree (Bhattacharyya et al., 2011). Random forests use separate training data between trees and randomly restrict the pool of attributes available when building each internal node (Bhattacharyya et al., 2011). Another method for reducing overfitting in decision trees is pruning, which involves removal of decision nodes without reducing the overall accuracy of the tree (Kirkos et al., 2007). These methods make random forests robust to overtraining and noise, and as each tree is generated independently there is little additional computational complexity (Bhattacharyya et al., 2011). Additionally, the only two parameters that require adjustment are the number of trees and the set of attributes to choose from when building each node, which makes decision forests simple to generate (Bhattacharyya et al., 2011).

4.7. Group method of data handling

Group method of data handling is an inductive data mining algorithm that calculates optimal solutions through a series of increasingly accurate models. It begins with a simple model, typically a polynomial of the form:

$$Z = w_0 + w_1x_1 + w_2x_2 + w_3x_1^2 + w_4x_2^2 + w_5x_1x_2 \tag{8}$$

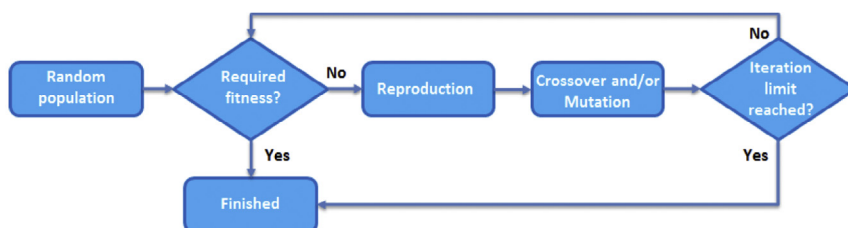


Fig. 7 – Standard genetic algorithm process.

After assessing the validity of the model further models are generated by using linear regression to calculate new coefficients for the polynomial. In this way group method of data handling is continually moving towards an optimal approximation. The generated models are functionally similar to a feed forward neural network, with the input to each neuron being derived from a polynomial representation of neurons in the preceding layer and a single neuron in the output layer. After each iteration the success of the model is calculated as the margin of error between actual and expected outputs. If this reaches a predetermined level of correctness the algorithm concludes (Ravisankar et al., 2011).

4.8. Text mining

Text mining refers to a specific form of data mining which is performed on plain text. Given this broad description there are a variety of different approaches to text mining, though many work by using one or more common pre-processing steps to transform the data into more quantitative samples:

- Filtering of stop words. Common joining words such as “the” or “is” are removed to simplify the dataset.
- Stemming. This involves reducing derived words to their common base, such as removing plurals and tenses.
- Identification of part-of-speech. This recognises the part of a sentence that a word occurs in and can be helpful in reducing error due to homophones and other language nuances.
- Word frequency analysis. The frequency of word i in document j is calculated as (tf_{ij}) .
- Word combination. Words are further combined into common concepts using a synonym list or more complicated measures.

Once the text has been converted to a more quantitative form a traditional data mining method is applied to actually perform the classification (Cecchini et al., 2010; Glancy and Yadav, 2011). An alternative to the above approach is to avoid pre-processing and deliberately include all raw data as a method of detecting abnormalities in the text. Measuring factors such as expression, complexity, and specificity within the text can be used to identify subtle differences between samples (Humpherys et al., 2011).

4.9. Self-organising map

Self-organising maps are a form of artificial neural network which consists of a single matrix of neurons. A non-linear algorithm is used to map inputs from a high-dimensional space to the two-dimensional array of neurons. The mapping is designed to model similar input vectors as neurons that are closer together in the resulting matrix, providing a visualisation of the inputs. A distance or neighbourhood function is used to group the nodes, such as the Euclidean distance formula or Gaussian formula (Quah and Sriganesh, 2008). The clustering function applied to each neuron is given by:

$$Y_{i+1} = Y_i + \alpha(X_i - Y_{i-1}) \quad (9)$$

where Y_i is the current weighting of a specific node, X_i is the current input vector, and α is the chosen distance function. The clustering step is performed a set number of iterations before the algorithm completes (Olszewski, 2014).

4.10. Process mining

Process mining involves analysis of transactions and event logs to construct models representing the behaviour of a system. A specific process instance is created to represent the individual cases within the system, and assumptions are applied to the available data to determine whether it is suitable for observation with this particular process. Given a model of expected behaviour within the system, the typical operation steps are:

- Log preparation and inspection. The logs are obtained and pre-processed to remove extraneous noise. If no existing model was available one can be created based on the provided logs.
- Analysis. Models are observed to check various behaviours. Aspects such as control flow, performance, and user-roles can be analysed to determine the expected outcomes of the system.
- Verification. Given the results of the above analysis, the process miner is applied to various samples and determines whether they represent typical system behaviour.

In addition to classification, the ability to generate its own model makes process mining useful for discovering flows and practices within a complicated system (Jans et al., 2011).

4.11. Artificial immune system

Artificial immune systems are a data mining approach that work by imitating the behaviour of a biological immune system to detect antigens (Wu and Banzhaf, 2008). A variety of biological characteristics can be simulated by the artificial immune system, but most model the creation of detector cells and their ability to detect foreign bodies. The detector cells are generated randomly and the simulation is performed to test and evaluate their effectiveness, similar to the training performed by other classification methods.

One common form is clonal selection, which continually generates detector cells that only live a short time. If a cell detects an antigen it extends its life to fight the intruder, and may also mutate as a result of the conflict. The surviving cells at the end of the simulation are therefore the ones best suited to detecting the antigens. Another common implementation is negative selection, which works by randomly creating cells and determining which of these react with other non-invasive cells within the system. Any that do are discarded, resulting in the remainder being proficient at detecting intruders (Soltani Halvaie and Akbari, 2014).

4.12. Hybrid methods

Hybrid methods are a combination of multiple traditional methods by selecting beneficial attributes of each in an attempt to create a superior algorithm for a specific problem domain.

Hybrid methods can be constructed in a variety of different ways: at the highest level methods can simply be applied linearly, with the outputs of the first providing the inputs of the second (Panigrahi et al., 2009). Similarly one method may be used as a pre-processing step to modify the data in preparation for classification (Jans et al., 2011), or at a lower level the individual steps of the algorithms can be intertwined to create something fundamentally original (Duman and Ozcelik, 2011). Additionally, hybrid methods can be used to tailor solutions to an individual problem domain. Different aspects of performance can be specifically targeted, including classification ability, ease of use, and computational efficiency.

5. Classification of existing financial fraud detection

In the following section we will classify the financial fraud detection techniques reviewed based on success rate, method chosen, and fraud type studied. This categorisation will enable us to demonstrate trends in current research methods, including which have been successful and any factors that have not been covered.

5.1. Classification based on performance

A variety of standards have been used to determine performance, but the three most commonly used are accuracy, sensitivity, and specificity. Accuracy measures the ratio of all successfully classified samples to unsuccessful ones. Sensitivity compares the amount of items correctly identified as fraud to the amount incorrectly listed as fraud, also known as the ratio of true positives to false positives. Specificity refers to the same concept with legitimate transactions, or the comparison of true negatives to false negatives (Bhattacharyya et al., 2011) (Ravisankar et al., 2011).

Tables 3, 4, and 5 classify financial fraud detection research based on these performance measures. Additionally, Fig. 9 depicts the amount of studies that have been performed on each combination of detection method and fraud type, and the comparative performance of each.

In addition to the three performance measures discussed here, several other performance measures have been used in the literature. For example, Duman et al. chose to show their results for sensitivity in graph form instead of deterministic values, grouped by each set of input parameters (Duman and Ozcelik, 2011). In addition to other forms of graphing (Quah and Sriganesh, 2008), some research used software-determined success levels or case-based procedures to determine the success of their fraud detection techniques (Sánchez et al., 2009; Jans et al., 2011).

From the results we can see that CI methods typically had better success rate than statistical methods. Sensitivity was slightly better for random forests and support vector machines than logistic regression, with comparable specificity and accuracy (Bhattacharyya et al., 2011). Genetic programming, support vector machines, probabilistic neural networks, and group method of data handling outperformed regression in all three areas (Ravisankar et al., 2011). Additionally, a neural

network with exhaustive pruning was found to be more specific and accurate than CDA (Bose and Wang, 2007). One statistical method seems to contradict this theory however: Bayesian belief networks were reported to be more accurate than neural networks and decision trees (Kirkos et al., 2007), and adding Bayesian logic to regression outperformed logistic regression alone (Bermúdez et al., 2008).

Most of the research showed a large difference between each method's sensitivity and specificity results. For example, Bhattacharyya et al. showed that logistic regression, support vector machines and random forests all performed significantly better at detecting legitimate transactions correctly than fraudulent ones (Bhattacharyya et al., 2011). Support vector machines, genetic programming, neural networks, group method of data handling, and particularly logistic regression were also slightly less sensitive (Ravisankar et al., 2011). And a neural network with exhaustive pruning showed more specificity than sensitivity (Bose and Wang, 2007).

As explained previously, fraud detection is a problem with a large difference in misclassification costs: it is typically far more expensive to misdiagnose a fraudulent transaction as legitimate than the reverse. With that in mind it would be beneficial for detection techniques to show a much higher sensitivity than specificity, meaning that these results are less than ideal. Contrary to this belief, Hoogs et al. hypothesised that financial statement fraud may carry higher costs for false positives, and their results reflect this with a much higher specificity (Hoogs et al., 2007). Panigrahi et al. also acknowledged the costs associated with following up credit card transactions marked as fraudulent, focussing their results on sensitivity only (Panigrahi et al., 2009). The CDA and CART methods, as well as the neural networks, Bayesian belief networks and decision trees performed better in this regard, with all showing a somewhat higher ability to classify fraudulent transactions than legitimate ones (Bose and Wang, 2007; Kirkos et al., 2007).

5.2. Classification based on detection algorithm

Classifying fraud detection practices by the detection algorithm is a useful way to identify the suitable techniques for this problem domain. It can also help us to determine why particular methods were chosen or successful. Additionally, we can identify any gaps in research by looking at algorithms which have not been explored sufficiently. Table 6 shows classification of financial fraud detection practices based on detection algorithm (conventional data mining and CI-based approaches) used.

Previously it was mentioned that early fraud detection research focussed on statistical models and neural networks; however, it may be noted that these methods still continue to be popular. Many used at least one form of neural network (Kirkos et al., 2007; Ravisankar et al., 2011; Bose and Wang, 2007), some investigated logistic regression (Bhattacharyya et al., 2011; Pinquet et al., 2007; Viaene et al., 2007; Ravisankar et al., 2011; Huang, 2013), while others applied Bayesian belief networks and other Bayesian algorithms (Holton, 2009; Kirkos et al., 2007; Bermúdez et al., 2008). Application of CDA has been relatively uncommon (Bose and Wang, 2007). Neural networks and logistic regression are often chosen for their well-established popularity, giving them the ability to be used as a control

Table 3 – Accuracy results.

Research	Fraud investigated	Method investigated	Accuracy
Bhattacharyya et al. (2011)	Credit card transaction fraud from a real world example	Logistic model (regression)	96.6–99.4%
		Support vector machines	95.5–99.6%
Olszewski (2014)	Credit card transaction fraud from a bank in Warsaw	Random forests	97.8–99.6%
		Self-organising map	100%
Soltani Halvaiee and Akbari (2014)	Credit card fraud from a Brazilian bank	Artificial immune system	94.6–96.4%
		Logistic regression	60.680%
Bermúdez et al. (2008)	Insurance fraud from automobile insurance claims for a Spanish company	Bayesian skewed regression	99.538%
		Decision trees	73.6%
Kirkos et al. (2007)	Financial statement fraud from a selection of Greek manufacturing firms	Neural networks	80%
		Bayesian belief networks	90.3%
Ravisankar et al. (2011)	Financial statement fraud with financial items from a selection of public Chinese companies	Support vector machine	70.41–73.41%
		Genetic programming	89.27–94.14%
Glancy and Yadav (2011)	Financial statement fraud with managerial statements for US companies	Neural network (feed forward)	75.32–78.77%
		Group method of data handling	88.14–93.00%
Cecchini et al. (2010)	Financial statement fraud with managerial statements for US companies	Logistic model (regression)	66.86–70.86%
		Neural network (probabilistic)	95.64–98.09%
Humpherys et al. (2011)	Financial statement fraud with managerial statements for US companies	Text mining with singular validation decomposition vector	95.65%
		Text mining	45.08–75.41%
Huang (2013)	Financial statement fraud from Taiwanese lawsuits	Text mining and support vector machine hybrid	50.00–81.97%
		Text mining and decision tree hybrid	67.3%
Bose and Wang (2007)	Financial statement fraud with financial items from a selection of public Chinese companies	Text mining and Bayesian belief network hybrid	67.3%
		Text mining and support vector machine hybrid	65.8%
Bose and Wang (2007)	Financial statement fraud with financial items from a selection of public Chinese companies	Logistic regression	19%–79%
		Support vector machine	71%–92%
Bose and Wang (2007)	Financial statement fraud with financial items from a selection of public Chinese companies	CDA	71.37%
		CART	72.38%
Bose and Wang (2007)	Financial statement fraud with financial items from a selection of public Chinese companies	Neural network (exhaustive pruning)	77.14%

method by which other techniques are tested. Comparatively, more advanced methods such as support vector machines and genetic programming have received substantially less attention. Yue et al. also reported that all the methods mentioned in their research were a form of classification, with no studies performed on clustering or time-series approaches, and that most of the research focussed on supervised learning as opposed to unsupervised (Yue et al., 2007).

Several of the research focussed on a single form of fraud detection which they advocated above others, such as studying text mining with the singular validation decomposition vector (Glancy and Yadav, 2011), self-organising maps (Quah and Sriganesh, 2008; Olszewski, 2014), logistic regression (Viaene et al., 2007; Pinquet et al., 2007), and fuzzy logic (Sánchez et al., 2009). Additionally, some researchers focussed solely on classification and regression trees (Bai et al., 2008; Sahin et al., 2013), Bayesian belief networks (Holton, 2009), individual statistical techniques (Panigrahi et al., 2009), artificial immune systems

(Wong et al., 2012; Soltani Halvaiee and Akbari, 2014), or their own hybrid methods (Duman and Ozcelik, 2011). This unilateral approach is useful for demonstrating the ability of the specific method in isolation, but without comparing it to other methods it is difficult to understand the relative performance of the technique. Additional factors such as the fraud type researched and the specific dataset used can influence the results of the experiment. Future research could focus on reviewing these methods as against other more established techniques.

A rising trend in fraud detection is the use of hybrid methods which utilise the strengths of multiple algorithms to classify samples. Duman and Ozcelik used a combination of scatter search and genetic algorithm, based on the latter but targeting attributes of scatter search such as the smaller populations and recombination as the reproduction method (Duman and Ozcelik, 2011). A different approach was taken by Panigrahi et al. who used two methods sequentially, beginning with the

Table 4 – Sensitivity results.

Research	Fraud investigated	Method investigated	Sensitivity
Bhattacharyya et al. (2011)	Credit card transaction fraud from a real world example	Logistic model (regression)	24.6–74.0%
		Support vector machines	43.0–68.7%
Olszewski (2014)	Credit card transaction fraud from a bank in Warsaw	Random forests	42.3–81.2%
		Self-organising map	100%
Soltani Halvaiee and Akbari (2014)	Credit card fraud from a Brazilian bank	Artificial immune system	33.6%–52.6%
Bermúdez et al. (2008)	Insurance fraud from automobile insurance claims for a Spanish company	Logistic regression	85.149%
		Bayesian skewed regression	85.149%
Kirkos et al. (2007)	Financial statement fraud from a selection of Greek manufacturing firms	Decision trees	75.0%
		Neural networks	82.5%
Ravisankar et al. (2011)	Financial statement fraud with financial items from a selection of public Chinese companies	Bayesian belief networks	91.7%
		Support vector machine	55.43–73.60%
		Genetic programming	85.64–95.09%
		Neural network (feed forward)	67.24–80.21%
		Group method of data handling	87.44–93.46%
Glancy and Yadav (2011)	Financial statement fraud with managerial statements	Logistic model (regression)	62.91–65.23%
		Neural network (probabilistic)	87.53–98.09%
		Text mining with singular validation decomposition vector	95.65%
Huang (2013)	Financial statement fraud from Taiwanese lawsuits	Logistic regression	24–93%
Bose and Wang (2007)	Financial statement fraud with financial items from a selection of public Chinese companies	Support vector machine	76–98%
		CDA	61.96%
Panigrahi et al. (2009)	Credit card fraud using legitimate customer transaction history as well as generic fraud transactions	CART	72.40%
		Neural network (exhaustive pruning)	80.83%
		Bayesian learning with Dempster-Shafer combination	71–83%
Hoogs et al. (2007)	Financial statement fraud from Accounting and Auditing Enforcement Releases by the Securities and Exchange Commission	Genetic algorithm	13–27%
Wu and Banzhaf (2008)	Transactional fraud in automated bank machines and point of sale from a financial institution	Coevolution artificial immune system	97.688–98.266%
		Standard evolution artificial immune system	92.486–95.376%

Depster-Schaefer method to combine rules and then using a Bayesian learner to detect the existence of fraud (Panigrahi et al., 2009). Some researchers applied fuzzy logic to introduce variation to their samples, attempting to transform it to resemble real world data before deploying a different technique to actually detect the presence of fraud (Jans et al., 2011). The investigators recognised that applying “fuzziness” to their problem increased the performance of their solution (Wu and Banzhaf, 2008). Similarly, several researchers combined traditional computational intelligence methods with text mining to analyse financial statements for the presence of fraud (Cecchini et al., 2010; Humpherys et al., 2011; Dong et al., 2014).

5.3. Classification based on fraud type

Given the varying nature of each type of fraud, the problem domain can differ significantly depending on the form that is being detected. By classifying the existing practices on the type of fraud investigated we can identify the techniques more suit-

able and more commonly used for a specific type of fraud. Additionally we can infer the varieties which are considered the most important for investigation depending on the scope and scale of their impact. Table 7 depicts the classification based on fraud types considered, along with the detection methods used. Table 8 lists the datasets used by each researcher.

As with each chosen algorithm, feature selection will differ depending on the problem domain. Specific financial statement fraud exists within individual companies, and as such attribute ratios are used instead of absolute values. Koh and Low provide a good example of the relevant ratios such as net income to total assets, interest payments to earnings before interest and tax, and market value of equity to total assets (Koh and Low, 2004). In comparison, research into credit card fraud has typically selected independent variables or aggregate values which may be quantitative or qualitative. For example, Bhattacharyya et al. made use of transaction amount, categorical values such as account number, transaction date, and currency, and aggregated properties like total transaction

Table 5 – Specificity results.

Research	Fraud investigated	Method investigated	Specificity
Bhattacharyya et al. (2011)	Credit card transaction fraud from a real world example	Logistic model (regression)	96.7–99.8%
		Support vector machines	95.7–99.8%
Soltani Halvaiee and Akbari (2014)	Credit card fraud from a Brazilian bank	Random forests	97.9–99.8%
		Artificial immune system	97.8–98.1%
Bermúdez et al. (2008)	Insurance fraud from automobile insurance claims for a Spanish company	Logistic regression	60.430%
Kirkos et al. (2007)	Financial statement fraud from a selection of Greek manufacturing firms	Bayesian skewed regression	99.677%
		Decision trees	72.5%
Ravisankar et al. (2011)	Financial statement fraud with financial items from a selection of public Chinese companies	Neural networks	77.5%
		Bayesian belief networks	88.9%
		Support vector machine	70.41–73.41%
		Genetic programming	89.27–94.14%
		Neural network (feed forward)	75.32–78.77%
Glancy and Yadav (2011)	Financial statement fraud with managerial statements	Group method of data handling	88.34–95.18%
		Logistic model (regression)	70.66–78.88%
Huang (2013)	Financial statement fraud from Taiwanese lawsuits	Neural network (probabilistic)	94.07–98.09%
		Text mining with singular validation decomposition vector	95.65%
Bose and Wang (2007)	Financial statement fraud with financial items from a selection of public Chinese companies	Logistic regression	24–91%
		Support vector machine	11–85%
Hoogs et al. (2007)	Financial statement fraud from Accounting and Auditing Enforcement Releases by the Securities and Exchange Commission	CDA	80.77%
		CART	72.36%
		Neural network (exhaustive pruning)	73.45%
Wu and Banzhaf (2008)	Transactional fraud in automated bank machines and point of sale from a financial institution	Genetic algorithm	98%–100%
		Coevolution artificial immune system	95.862–97.122%
		Standard evolution artificial immune system	99.311%

amount per day, and average amount spent at a single merchant (Bhattacharyya et al., 2011).

We can see that the existing research has been greatly unbalanced in fraud type studied. The vast majority of papers have focussed on two forms of financial fraud: *credit card fraud* and *financial statement fraud*. Only a handful of studies have looked at securities and commodities fraud, and many focussed on external forms of corporate fraud while neglecting the internal ones (Jans et al., 2011). Ngai et al. found that insurance fraud had received the highest coverage during their research (Ngai et al., 2011): the fact that we identified only a few examples of published literature on this type of fraud since 2007 indicates that research into insurance fraud is declining. Additionally, no studies have been performed directly on mortgage fraud or money laundering. Reasons for this disparity may include differing relevance to stakeholders of each fraud type.

To determine which method was the most successful for each fraud type we looked at the accuracy scores reported in the experimental research. There are many metrics used to assess performance but accuracy was the most commonly used and therefore provided the best basis for comparison. Analysis of the individual problem domain is required to completely assess the usefulness of each method, but the results shown

here should provide a useful starting point for future researchers to investigate new fraud detection algorithms. It should be noted that we have reported only the best accuracy result obtained by each method for specific fraud types; the test conditions for individual experiments are not necessarily comparable. Table 9 shows these results for each fraud type, as well as comparable methods.

All the solutions used for credit card fraud had a very high success rate, including regression, support vector machines, artificial immune systems, and random forests. Out of these the self organising map is recommended due to its perfect accuracy with a 10,000 transaction sample (Olszewski, 2014). As mentioned earlier there has been a lack of recent research into insurance fraud, with only one researcher comparing typical logistic regression to a hybrid version with Bayesian logic. In this case the hybrid method offered a significant improvement with an accuracy of 99.5% compared to 60.7% (Bermúdez et al., 2008). Financial statement fraud had a large variance in results, from CDA at 71.4% to a probabilistic neural network with 98.1% (Bose and Wang, 2007; Ravisankar et al., 2011). Several other methods also had success rate greater than 90%, including Bayesian belief networks, support vector machine, genetic programming, group method of data handling, and some hybrid methods based on text mining.

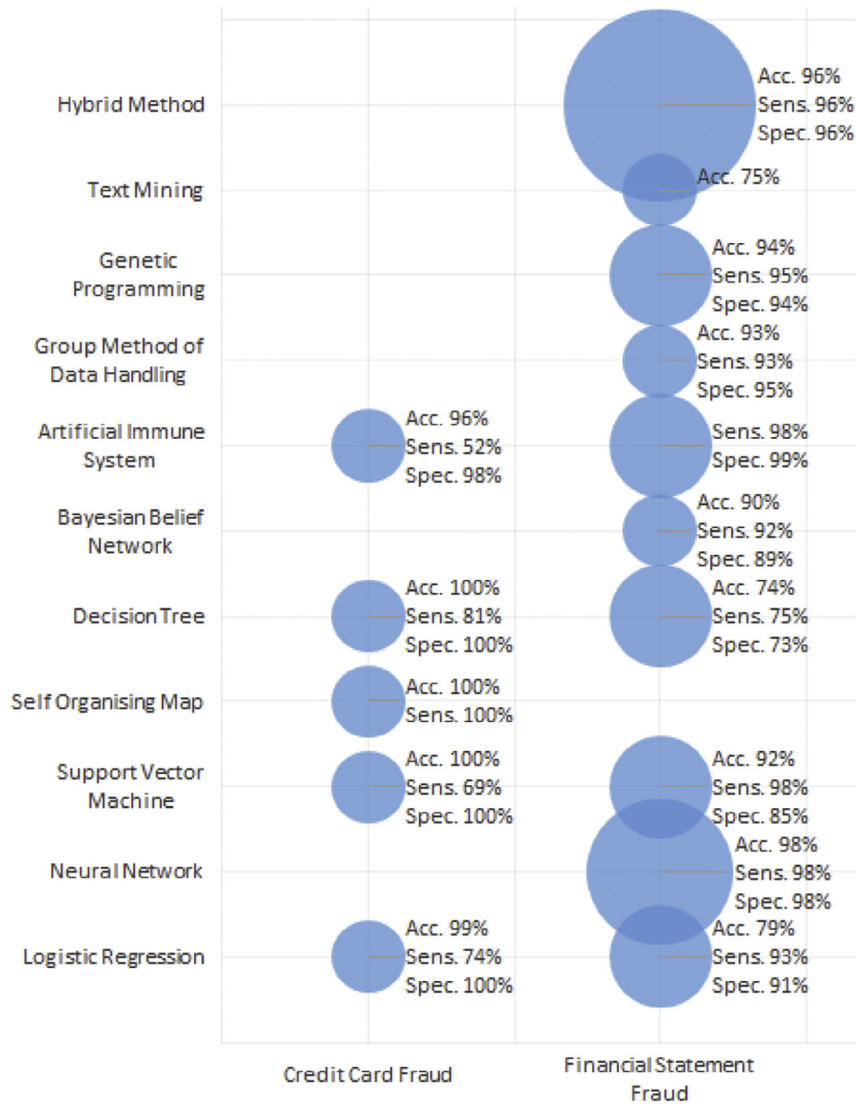


Fig. 9 – Comparative performance of various detection methods.

6. Financial fraud detection challenges and future directions

Financial fraud detection is an evolving field in which it is desirable to stay ahead of the perpetrators. Additionally, it is evident that there are still facets of intelligent fraud detection that have not been investigated. In this section we present some of the key issues associated with financial fraud detection and suggest areas for future research. Some of the identified issues and challenges are as follows:

- *Typical classification problems:* CI and data mining-based financial fraud detection is subject to the same issues as other classification problems, such as feature selection, parameter tuning, and analysis of the problem domain.
- *Fraud types and detection methods:* Financial fraud is a diverse field and there has been a large imbalance in both fraud types and detection methods studied: some have been studied extensively while others, such as hybrid methods, have only been looked superficially.
- *Privacy considerations:* Financial fraud is a sensitive topic and stakeholders are reluctant to share information on the subject. This has led to experimental issues such as undersampling.
- *Computational performance:* As a high-cost problem it is desirable for financial fraud to be detected immediately. Very little research has been conducted on the computational performance of fraud detection methods for use in real-time situations.
- *Evolving problem:* Fraudsters are continually modifying their techniques to remain undetected. As such detection methods are required to be able to constantly adapt to new fraud techniques.
- *Disproportionate misclassification costs:* Fraud detection is primarily a classification problem with a vast difference in misclassification costs. Research on the performance of

Table 6 – Qualitative analysis of methods researched in existing fraud detection literature.

Method investigated	Fraud investigated	Research
Neural network	Financial statement fraud	Bose and Wang (2007); Kirkos et al. (2007); Ravisankar et al. (2011)
Logistic model	Credit card fraud Insurance fraud	Bhattacharyya et al. (2011) Pinquet et al. (2007); Viaene et al. (2007); Bermúdez et al. (2008)
Support vector machine	Financial statement fraud Credit card fraud	Ravisankar et al. (2011); Huang (2013) Bhattacharyya et al. (2011); Whitrow et al. (2009)
Decision trees, forests and CART	Financial statement fraud Credit card fraud	Ravisankar et al. (2011); Huang (2013) Bhattacharyya et al. (2011); Whitrow et al. (2009); Sahin et al. (2013)
Genetic algorithm/programming	Financial statement fraud	Bose and Wang (2007); Bai et al. (2008); Kirkos et al. (2007)
Text mining	Financial statement fraud	Ravisankar et al. (2011); Hoogs et al. (2007)
Group method of data handling	Financial statement fraud	Glancy and Yadav (2011); Cecchini et al. (2010); Zaki and Theodoulidis (2013)
Response-surface methodology	Financial statement fraud	Ravisankar et al. (2011)
Self-organising map	Credit card fraud	Zhang and Zhou (2004) Quah and Sriganesh (2008); Sánchez et al. (2009); Olszewski (2014)
Bayesian belief network	Corporate fraud	Holton (2009)
Process mining	Financial statement fraud	Kirkos et al. (2007)
Artificial immune system	Securities and commodities fraud Credit card fraud	Jans et al. (2011) Wu and Banzhaf (2008); Wong et al. (2012); Soltani Halvaie and Akbari (2014)
Hybrid methods	Credit card fraud Insurance fraud Financial statement fraud	Panigrahi et al. (2009); Duman and Ozcelik (2011) Bermúdez et al. (2008) Cecchini et al. (2010); Dong et al. (2014); Humpherys et al. (2011)
All/generic	All/generic	Yue et al. (2007); Ngai et al. (2011)

detection methods with respect to this factor is an area which needs further attention.

- *Generic framework*: Given that there are many varieties of fraud, a generic framework which can be applied to multiple fraud categories would be valuable.

As a classification problem, financial fraud detection suffers from the same issues as other similar problems. Feature selection has a high impact on the success of any classification method. While some researchers have mentioned feature selection for one type of fraud (Koh and Low, 2004; Bhattacharyya et al., 2011), no comparisons have been made between features for differing problem domains. Also, one of the major benefits of the computational intelligence and data mining methods is their ability to be adjusted to fit the problem domain. Existing research has rarely used any form of customisation or tuning for specific problems; however, tuning is an important factor in the context of an algorithm's performance. For example, the number of nodes and internal layers within a neural network has a large impact on both accuracy and computational performance. Similarly the kernel function chosen will considerably alter the success of a support vector machine and parameters such as the fitness function, crossover method, and probability for mutation will impact the results of a genetic programming algorithm. Research on customisation or tuning of the computational methods is required to truly comprehend the ability of each method. Further, in other data mining cases the solution algorithm is selected based on its performance within the problem domain, which for financial fraud detection is the type of fraud investigated. Studies on the suitability of various methods for each fraud category are necessary

to understand which attributes of each algorithm make them appropriate for detecting financial fraud.

From the existing literature it is apparent that there are some forms of fraud that have not been investigated as extensively as others. Financial statement fraud has been considerably investigated, which is understandable given its high profile nature, but there are other forms of fraud that have a significant impact on consumers. Credit card fraud often has a direct impact on the public and the recent increase in online transactions has led to a majority of the US public being concerned with identity theft (Bhattacharyya et al., 2011). A benefit of this close relation to the user is that credit card fraud is typically detected quickly, which gives researchers access to large datasets of unambiguous transactions. Other forms of fraud which have not been covered in depth include money laundering, mortgage, and securities and commodities fraud. A lack of sufficient sample size may be the reason for the lack of research in these areas (Ngai et al., 2011). Future studies that focussed on these types of fraud detection would be beneficial.

The private nature of financial data has led to institutions being reluctant to share fraudulent information. This has had an effect both on the fraud types that have been investigated as well as the datasets used for the purpose. In the published literature many of the financial fraud simulations consisted of less than a few hundred samples, typically with comparable amounts of fraudulent and legitimate specimens. This is contrary to the realities of the problem domain, where fraud cases are far outweighed by legitimate transactions (Bhattacharyya et al., 2011). Undersampling the problem domain like this can cause biases in the data that do not accurately represent real-world scenarios (Hoogs et al., 2007). There is a definite need

Table 7 – Methods studied by type of fraud investigated.

Fraud type	Method applied	Research on the type of fraud
Credit card	Support vector machines Decision trees Self-organising maps Fuzzy logic Artificial immune system Hybrid methods	Bhattacharyya et al. (2011) and Paasch (2010) investigated credit card fraud from an international operation. Quah and Sriganesh (2008) investigated a banking database from the Singapore branch of a well-known international bank. Sánchez et al. (2009) investigated fraud in multinational department stores. Duman and Ozcelik (2011) investigated typical consumer spending to determine fraud in a major bank in Turkey. Panigrahi et al. (2009) investigated variation in legitimate customer transaction behaviour with synthesised credit card data. Wu and Banzhaf (2008) investigated automated bank machines and point of sale from an anonymous financial institution. Whitrow et al. (2009) investigated credit card transactions from two separate banks. Wong et al. (2012) investigated transactions from a major Australian bank. Sahin et al. (2013) investigated six months of transactions from an anonymous bank. Olszewski (2014) investigated credit card accounts of residents in Warsaw. Soltani Halvaeie and Akbari (2014) investigated credit card transactions from a Brazilian bank.
Securities and commodities and other corporate	Bayesian belief network Process mining	Jans et al. (2011) investigated internal transactional fraud from a successful, anonymous European financial institution. Holton, 2009) Investigated emails and discussion group messages to detect corporate fraud.
Insurance fraud	Logistic model Hybrid methods	Pinquet et al. (2007) , Viaene et al. (2007) , and Bermúdez et al., 2008) all investigated motor insurance claims from Spanish insurance companies.
Financial statement	Response-surface methodology Neural networks Decision trees Bayesian belief networks Support vector machine Genetic algorithms Group method of data handling Logistic model (regression) Text mining Hybrid methods	Zhang and Zhou (2004) investigated financial statement fraud in general. Kirkos et al. (2007) investigated a selection of Greek manufacturing firms. Ravisankar et al. (2011) , Bose and Wang, 2007), and Bai et al., 2008) investigated a series of public Chinese companies. Glancy and Yadav (2011) and Humpherys et al., 2011) investigated managerial statements from official company documents. Hoogs et al. (2007) and Cecchini et al. (2010) investigated Accounting and Auditing Enforcement Releases authored by a selection of US companies, and Dong et al. (2014) investigated the same documents with a focus on Chinese companies. Zaki and Theodoulidis (2013) investigated the litigation section of the Securities Exchange Commission website. Huang (2013) investigated Taiwanese companies that had been accused of fraud.

for further studies with realistic samples to accurately depict the performance of each method ([Glancy and Yadav, 2011](#)).

Some forms of financial fraud occur very rapidly, such as credit card fraud. If a fraudster obtains an individual's credit card information it is very likely that they will use it immediately until the card limit is reached. The ability to detect fraud in real-time would be highly beneficial as it may be able to prevent the fraudster from making subsequent transactions. Computational performance is therefore a key factor to consider in fraud detection. Though some researchers have noted the performance of their particular methods ([Bhattacharyya et al., 2011](#); [Quah and Sriganesh, 2008](#)), most studies were simulations performed on test datasets. Further research focussing on the computational as well as classification performance is required.

Unlike many classification problems, fraud detection solutions must be capable of handling active attempts to circumvent them. As detection methods become more intelligent, fraudsters are also constantly upgrading their techniques. For example, in the last few decades credit card fraud has moved from individuals stealing or forging single cards to large-scale phone and online fraud perpetrated by organised groups

([Bhattacharyya et al., 2011](#)). It is therefore necessary for fraud detection methods to be capable of evolving to stay ahead of fraudsters. Some researchers have considered models for adaptive classification, however further research is required to fully develop these for use in practical fraud detection problems ([Zhou and Kapoor, 2011](#)).

As explained previously fraud has a large cost to businesses. Additionally, fraud detection has associated costs: systems require maintenance and computational power, and auditors must be employed to monitor them and investigate when a potential fraud case is identified ([Kirkos et al., 2007](#)). The expense of a false positive, in misclassifying a legitimate transaction as fraud, is typically far less than that of a false negative ([Ngai et al., 2011](#)). Insufficient study has been performed on the disproportionate nature of these costs, with attention typically focussing on the traditional classification performance methods outlined in [Section 5.1](#). Considering the accuracy of each fraud detection method, focus should be on achieving an optimum balance for each technique such that the expense is smallest. Research specifically focussed on finding this balance would add significant real-world value to financial fraud detection.

Table 8 – Datasets used in fraud detection research.

Fraud type	Research	Dataset
Credit card	Bhattacharyya et al. (2011); Paasch (2010), Quah and Sriganesh (2008)	Almost 50 million credit card transactions from an international company from 2006 to 2007, all of which occurred in a single country. Approximately 100 transactions per customer for a single month at the Singaporean branch of an anonymous, well-known bank.
	Sánchez et al. (2009)	Credit card transactions provided by a selection of retail companies in Chile, of which approximately 0.3% were fraudulent.
	Duman and Ozcelik (2011)	Approximately 250,000 credit card transactions from an anonymous Turkish bank, including 1050 that are fraudulent.
	Panigrahi et al. (2009)	Unknown amount of synthesised credit card transactional data.
	Wu and Banzhaf (2008)	Selection of 522,728 legitimate and 346 fraudulent credit card transactions from an anonymous financial institution.
	(Whitrow et al. (2009)	Credit card transactions from two banks during 2005, including 176 million transactions from nearly 17 million accounts of which 14,281 experienced fraud.
	Wong et al. (2012)	Total of 640,361 transactions for 21,746 credit cards from a major Australian bank
	Sahin et al. (2013)	Deliberate undersampling of approximately 11,344,000 credit card transactions over a 6 month period, 484 of which were legitimate.
	Olszewski (2014)	10000 credit card accounts from a bank in Warsaw from January to March 2005, of which 100 were fraudulent.
	Soltani Halvaeie and Akbari (2014)	Anonymous transaction records from a Brazilian bank, of which 3.74% were fraudulent in some way.
Securities and commodities and other corporate	Jans et al. (2011)	Random selection of 10,000 process instances from an anonymous European financial institution in 2007.
	Holton (2009)	Random sample of 50 disgruntled and 40 non-disgruntled message obtained from electronic, intra-company communications.
Insurance fraud	Pinquet et al. (2007)	Selection of 4970 audited motor insurance claims from 2000, including 681 cases of fraud.
	Viaene et al. (2007)	Selection of 2403 audited motor insurance claims from 2000, including 174 cases of fraud.
	Bermúdez et al. (2008)	Selection of 10,000 audited motor insurance claims from 2000, including 101 cases of fraud.
Financial statement	Kirkos et al. (2007)	Financial statements from 76 Greek manufacturing firms, of which 38 had demonstrated fraud.
	Ravisankar et al. (2011)	Financial statements from 202 companies listed on Chinese stock exchanges, of which 101 had demonstrated fraud.
	Bose and Wang (2007)	Financial statements from a selection of 202 Chinese companies.
	Bai et al. (2008)	Ten known Chinese companies that have been found guilty of fraud between 1999 and 2002
	Glancy and Yadav, 2011)	10-K documents from 2006 to 2008 for various known-fraudulent US companies, matched with a selection of similar, legitimate companies.
	Humpherys et al. (2011)	202 10-K documents for US companies from 1995–2004, of which half were fraudulent.
	Hoogs et al. (2007)	Observations for 390 companies from US Accounting and Auditing Enforcement Releases between 1991 and 2004, of which 51 companies demonstrated fraud.
	Cecchini et al. (2010)	10-K documents from 61 fraudulent and 61 non-fraudulent US companies from between 1993 and 2002.
	Zaki and Theodoulidis (2013)	Sixty-two litigation releases from the Securities and Exchange Commission website.
	Huang (2013)	A series of companies prosecuted for fraud in Taiwan between 1998 and 2009, matched with similar legitimate companies in a ratio of 1:4.
Dong et al. (2014)	10-K documents from 17 fraudulent and 17 non fraudulent Chinese companies from 2003 to 2014.	

Given the diversity of common categories of fraud it would be useful to have some form of generic framework that could apply to more than one fraud category. Such a framework could be used to study the differences between various types of fraud, or even specific details such as differentiating between stolen and counterfeit credit cards (Bhattacharyya et al., 2011). A ubiquitous model could also be used to determine which specific fraud detection method is applicable given the problem domain. This approach has been investigated slightly with response surface methodology (Zhou and Kapoor, 2011), but more detailed research is desirable.

7. Conclusion

Fraud detection is an important part of the modern finance industry. This literature review studied research into intelligent approaches to fraud detection, both statistical and computational. Though their performance differed, each technique was shown to be reasonably capable at detecting various forms of financial fraud. In particular, the ability of the computational methods such as neural networks and support vector machines to learn and adapt to new techniques is highly effective to the evolving tactics of fraudsters.

Table 9 – Most successful methods for each fraud type based on the accuracy measure.

Fraud investigated	Best performer		Comparative Performers	
	Method	Accuracy	Method	Accuracy
Credit card fraud	Self-organising map	100%	Logistic regression	99.4%
			Support vector machine	99.6%
			Random forest	99.6%
			Artificial immune system	96.4%
Insurance fraud	Hybrid method	99.5%	Logistic regression	60.7%
Financial statement fraud	Neural network	98.1%	Decision tree	73.6%
			Bayesian belief network	90.3%
			Support vector machine	92.0%
			Genetic programming	94.1%
			Group method of data handling	93.0%
			Logistic regression	79.0%
			Hybrid method	95.7%
			Text mining	75.4%
			GDA	71.4%

There are still many aspects of intelligent fraud detection that have not yet been the subject of research. Some types of fraud, as well as some data mining methods, have been superficially explored but require future study to be completely understood. There is also the opportunity to examine the performance of existing methods by adjusting their parameters, as well as the potential to study cost benefit analysis of computational fraud detection. Finally, further research into the differences between each type of financial fraud could lead to a general framework which would greatly improve the accuracy of intelligent detection methods.

REFERENCES

- Abbott LJ, Park Y, Parker S. The effects of audit committee activity and independence on corporate fraud. *Managerial Finance* 2000;26:55–68.
- Aleskerov E, Freisleben B, Rao B. Cardwatch: a neural network based database mining system for credit card fraud detection. In *Computational Intelligence for Financial Engineering (CIFER)*, 1997, Proceedings of the IEEE/IAFE 1997. pp. 220–6, IEEE: New York. 1997.
- Bai B, Yen J, Yang X. False financial statements: characteristics of China's listed companies and CART detecting approach. *International Journal of Information Technology & Decision Making* 2008;7:339–59.
- Bell TB, Carcello JV. A decision aid for assessing the likelihood of fraudulent financial reporting. *Auditing: A Journal of Practice & Theory* 2000;19:169–84.
- Beneish MD. The detection of earnings manipulation. *Financial Analysts Journal* 1999;55:24–36.
- Bermúdez L, Pérez J, Ayuso M, Gómez E, Vázquez F. A Bayesian dichotomous model with asymmetric link for fraud in insurance. *Insur Math Econ* 2008;42:779–86.
- Bhattacharyya S, Jha S, Tharakunnel K, Westland JC. Data mining for credit card fraud: a comparative study. *Decision Support Systems* 2011;50:602–13.
- Bolton RJ, Hand DJ. Unsupervised profiling methods for fraud detection. *Credit Scoring and Credit Control VII* 2001;5–7.
- Bolton RJ, Hand DJ. Statistical fraud detection: a review. *Statistical Science* 2002;17:235–55.
- Bose I, Wang J. Data mining for detection of financial statement fraud in Chinese Companies. Paper presented at the International conference on electronic commerce, administration, society and education, Hong Kong, 15–17 August 2007. 2007.
- Busta B, Weinberg R. Using Benford's Law and neural networks as a review procedure. *Managerial Auditing Journal* 1998;13:356–66.
- Calderon TG, Cheh JJ. A roadmap for future neural networks research in auditing and risk assessment. *International Journal of Accounting Information Systems* 2002;3:203–36.
- Cecchini M, Aytug H, Koehler GJ, Pathak P. Making words work: Using financial text as a predictor of financial events. *Decision Support Systems* 2010;50:164–75.
- Deshmukh A, Talluru L. A rule-based fuzzy reasoning system for assessing the risk of management fraud. *International Journal of Intelligent Systems in Accounting, Finance & Management* 1998;7:223–41.
- Dong W, Liao SS, Fang B, Cheng X, Chen Z, Fan W. The detection of fraudulent financial statements: an integrated language model. *PACIS 2014 proceeding*, Paper 383. 2014.
- Duman E, Ozcelik MH. Detecting credit card fraud by genetic algorithm and scatter search. *Expert Systems with Applications* 2011;38:13057–63.
- Edge ME, Sampaio PRF. A survey of signature based methods for financial fraud detection. *Computers & Security* 2009;28(6):381–94.
- Fanning KM, Cogger KO. Neural network detection of management fraud using published financial data. *International Journal of Intelligent Systems in Accounting, Finance & Management* 1998;7:21–41.
- FBI, Federal Bureau of Investigation, Financial Crimes Report to the Public Fiscal Year, Department of Justice, United States. 2010–2011. Available from: <http://www.fbi.gov/stats-services/publications/financial-crimes-report-2010-2011>.
- Feroz EH, Kwon TM, Pastena VS, Park K. The efficacy of red flags in predicting the SECs targets: an artificial neural networks approach. *International Journal of Intelligent Systems in Accounting, Finance & Management* 2000;9:145–57.
- Fraser IA, Hatherly DJ, Lin KZ. An empirical investigation of the use of analytical review by external auditors. *The British Accounting Review* 1997;29:35–47.
- Glancy FH, Yadav SB. A computational model for financial reporting fraud detection. *Decision Support Systems* 2011;50:595–601.
- Green BP, Choi JH. Assessing the risk of management fraud through neural network technology. *Auditing-A Journal Of Practice & Theory* 1997;16:14–28.

- Hansen J, McDonald JB, Messier Jr W, Bell TB. A generalized qualitative-response model and the analysis of management fraud. *Manage Sci* 1996;42:1022–32.
- Holton C. Identifying disgruntled employee systems fraud risk through text mining: a simple solution for a multi-billion dollar problem. *Decision Support Systems* 2009;46:853–64.
- Hoogs B, Kiehl T, Lacombe C, Senturk D. A genetic algorithm approach to detecting temporal patterns indicative of financial statement fraud. *Intelligent Systems in Accounting, Finance and Management* 2007;15:41–56.
- Huang SY. Fraud detection model by using support vector machine techniques. *International Journal of Digital Content Technology and its Applications* 2013;7:32–42.
- Humpherys SL, Moffitt KC, Burns MB, Burgoon JK, Felix WF. Identification of fraudulent financial statements using linguistic credibility analysis. *Decision Support Systems* 2011;50:585–94.
- Jans M, van der Werf JM, Lybaert N, Vanhoof K. A business process mining application for internal transaction fraud mitigation. *Expert Systems with Applications* 2011;38:13351–9.
- Kirkos E, Spathis C, Manolopoulos Y. Data mining techniques for the detection of fraudulent financial statements. *Expert Systems with Applications* 2007;32:995–1003.
- Koh HC, Low CK. Going concern prediction using data mining techniques. *Managerial Auditing Journal* 2004;19:462–76.
- Kou Y, Lu C-T, Sirwongwattana S, Huang Y-P. Survey of fraud detection techniques. In: 2004 IEEE international conference on networking, sensing and control, vol. 2. New York: IEEE; 2004. p. 749–54.
- Maes S, Tuyls K, Vanschoenwinkel B, Manderick B. Credit card fraud detection using Bayesian and neural networks. In *Proceedings of the 1st International Naiso Congress on Neuro Fuzzy Technologies*. 2002.
- Ngai E, Hu Y, Wong Y, Chen Y, Sun X. The application of data mining techniques in financial fraud detection: a classification framework and an academic review of literature. *Decision Support Systems* 2011;50:559–69.
- Olszewski D. Fraud detection using self-organizing map visualizing the user profiles. *Knowledge-Based Systems* 2014;70:324–34.
- Paasch CA. In *Credit card fraud detection using artificial neural networks tuned by genetic algorithms*. Hong Kong Univ. of Sci. and Tech. (Hong Kong). 2010.
- Panigrahi S, Kundu A, Sural S, Majumdar AK. Credit card fraud detection: a fusion approach using Dempster–Shafer theory and Bayesian learning. *Information Fusion* 2009;10:354–63.
- Phua C, Lee V, Smith K, Gayler R. A comprehensive survey of data mining-based fraud detection research. *arXiv preprint arXiv:1009.6119*. 2010.
- Pinquet J, Ayuso M, Guillen M. Selection bias and auditing policies for insurance claims. *Journal of Risk and Insurance* 2007;74:425–40.
- Qah JT, Sriganesh M. Real-time credit card fraud detection using computational intelligence. *Expert Systems with Applications* 2008;35:1721–32.
- Ravisankar P, Ravi V, Raghava Rao G, Bose I. Detection of financial statement fraud and feature selection using data mining techniques. *Decision Support Systems* 2011;50:491–500.
- Rezaee Z. In *Financial statement fraud: prevention and detection*. Hoboken, NJ: John Wiley & Sons; 2002.
- Sánchez D, Vila M, Cerda L, Serrano J-M. Association rules applied to credit card fraud detection. *Expert Systems with Applications* 2009;36:3630–40.
- Sahin Y, Bulkan S, Duman E. A cost-sensitive decision tree approach for fraud detection. *Expert Systems with Applications* 2013;40:5916–23.
- Sohl JE, Venkatachalam A. A neural network approach to forecasting model selection. *Information & Management* 1995;29:297–303.
- Soltani Halvaeie N, Akbari MK. A novel model for credit card fraud detection using artificial immune systems. *Applied Soft Computing* 2014;24:40–9.
- Spathis C, Doumpos M, Zopounidis C. Detecting falsified financial statements: a comparative study using multicriteria analysis and multivariate statistical techniques. *European Accounting Review* 2002;11: 509–35.
- Spathis CT. Detecting false financial statements using published data: some evidence from Greece. *Managerial Auditing Journal* 2002;17:179–91.
- Stolfo S, Fan W, Lee W, Prodromidis A, Chan P. Credit card fraud detection using meta-learning: Issues and initial results. In *AAAI-97 workshop on fraud detection and risk management*. 1997.
- Summers SL, Sweeney JT. Fraudulently misstated financial statements and insider trading: an empirical analysis. *Accounting Review* 1998;73:131–46.
- Vatsa V, Sural S, Majumdar AK. A game-theoretic approach to credit card fraud detection. In: *Information systems security*. Berlin, Germany: Springer; 2005. p. 263–76.
- Viaene S, Ayuso M, Guillen M, Van Gheel D, Dedene G. Strategies for detecting fraudulent claims in the automobile insurance industry. *Eur J Oper Res* 2007;176:565–83.
- West J, Bhattacharya M, Islam R. “Intelligent financial fraud detection practices: an investigation”, in *Proceedings of the 10th international conference on security and privacy in communication networks (SecureComm 2014)*. 2014.
- Whitrow C, Hand DJ, Juszcak P, Weston D, Adams NM. Transaction aggregation as a strategy for credit card fraud detection. *Data Mining and Knowledge Discovery* 2009;18:30–55.
- Wong N, Ray P, Stephens G, Lewis L. Artificial immune systems for the detection of credit card fraud: an architecture, prototype and preliminary results. *Information Systems Journal* 2012;22:53–76.
- Wu SX, Banzhaf W. Combatting financial fraud: a coevolutionary anomaly detection approach. In *proceedings of the 10th annual conference on genetic and evolutionary computation*. pp. 1673–80, ACM. 2008.
- Yang W-S, Hwang S-Y. A process-mining framework for the detection of healthcare fraud and abuse. *Expert Systems with Applications* 2006;31:56–68.
- Yeh I, Lien C-H. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications* 2009;36:2473–80.
- Yue D, Wu X, Wang Y, Li Y, Chu C-H. A review of data mining-based financial fraud detection research. In *WiCom 2007 international conference on Wireless Communications, Networking and Mobile Computing*, 2007. pp. 5519–22, IEEE. 2007.
- Zaki M, Theodoulidis B. Analyzing financial fraud cases using a linguistics-based text mining approach. Available at SSRN 2353834. 2013.
- Zhang D, Zhou L. Discovering golden nuggets: data mining in financial application. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* 2004;34:513–22.
- Zhang G, Eddy Patuwo B, Hu MV. Forecasting with artificial neural networks: the state of the art. *Int J Forecast* 1998;14:35–62.
- Zhou W, Kapoor G. Detecting evolutionary financial statement fraud. *Decision Support Systems* 2011;50:570–5.

Jarrold West is currently studying Honours in the School of Computing & Mathematics, Charles Sturt University, Australia. Jarrod completed his Bachelor degree in Computer Science from the University of Newcastle, Australia in 2013. As a researcher he is primarily interested in various aspects of information security. Jarrod's current research is focussed on computational intelligence and data mining based financial fraud detection.

Maumita Bhattacharya currently works as an academic in the School of Computing & Mathematics, Charles Sturt University, Australia.

As an academic Maumita has been involved in research and teaching in various disciplines of computer science and engineering including artificial intelligence, data mining, information security and computer ethics. Her current research interests focus on design and development of advanced computational intelligence algorithms and methods and their applications to information security, data mining, optimisation and other complex real world problem domains. Maumita serves in the editorial boards and scientific committees of many reputed journal and conferences in these fields.