Streamlining the Detection of Accounting Fraud through Web Mining and Interpretable Internal Representations

Duarte Trigueiros
University of Macau, University Institute of Lisbon
Lisbon, Portugal
Email: dmt@iscte.pt

Carolina Sam
Master of European Studies Alumni Association
Macau, China
Email: kasm@customs.gov.mo

***Abstract*—Considerable effort has been devoted to the development of Artificial Intelligence tools able to support the detection of fraudulent accounting reports. Published results are promising but, till the present date, the use of such research has been limited, due to the "black box" character of the developed tools and the cumbersome input task they require. The tool described in this paper solves both problems while improving specificity of diagnostics. It is based on Web Mining and on Multilayer Perceptron classifiers where a modified learning method leads to meaningful representations. Such representations are then input to a features' map where trajectories towards or away from fraud and other features are identified. The final result is a robust Web Mining-based, self-explanatory fraud detection tool.**

Keywords–Type of Information Mining; Knowledge Extraction; Accounting Fraud Mining.

I. INTRODUCTION

Fraud may cost US companies over USD 400 billion annually [1]. Amongst different types of fraud, manipulation of accounting reports is paramount. In spite of measures put in place to detect fraudulent book-keeping, manipulation is still ongoing, probably on a huge scale [1]. Auditors are required to assess the plausibility of manipulated reports. They apply analytical procedures to inspect sets of transactions which are the building blocks of reports. But detecting fraud internally is a difficult task as managers deliberately try to deceive auditors. Most material frauds stem from the top levels of the organization where controls are least effective. The general belief is that analytic procedures alone are rarely effective in detecting fraud [2].

In response to concerns about audit effectiveness in detecting fraud, quantitative techniques are being applied to the modelling of relationships underlying published reports' data with a view to discriminate between fraudulent and non-fraudulent reports [3][4]. Such external, *ex-post* approach would be valuable as a tool in the hands of users of published reports such as investors, analysts and banks. Artificial Intelligence (AI) techniques are likewise being developed to the same end. Detailed review articles covering this research are available [5][6].

A discouraging fact is that analysts do not use AI tools designed to help detecting accounting manipulation. This is largely due to the fact that such tools are "black boxes" where results cannot be explained using the viewpoint, language and expertize of analysts [2]. Since analysts are responsible for their decisions, tools they may use to support decisions must be transparent and self-explanatory. Moreover, extract, transform and load (ETL) tasks required by such AI tools are time-consuming and difficult to automate in this case. The paper describes work-in-progress seeking to overcome the above limitations. Web Mining is first employed to find, download and store accounting data. Then, fraud and two other attributes known to widen fraud propensity space are predicted by three Multilayer Perceptron (MLP) classifiers where a modified learning method leads to internal representations similar to financial ratios, readily interpretable by analysts. Such ratios then input a features' map where trajectories towards or away from fraud and other features are visualized. Diagnostic interpretation is further enhanced with the display of cases similar to those being analyzed.

The objective of the tool is not so much to innovate but to streamline a well-known but opaque and cumbersome practice. Its sole original contribution is the strict adherence to users requirements including a new MLP training method leading to transparent diagnostic. Fraud detection covers many types of deception: plagiarism, credit card fraud, medical prescription fraud, false insurance claim, insider trading, accounting reports' manipulation and other [12][13]. Frameworks used in the detection of, say, credit card fraud (such as Game Theory), are not necessarily efficient in detecting other types of deception. Neural Networks are widely used in research devoted to the detection of accounting fraud [7][8][9][10][11] and reported performance is satisfactory.

Section II describes data and models while offering extensive methodological details. Section III reports preliminary results and presents the architecture of the tool to be deployed. Section IV discusses expected benefits.

II. METHODOLOGY

A. Accounting Information

An accounting report is a collection of monetary amounts with an attached meaning: revenues of the period, different types of expenses, asset values at the end of the period, liabilities and others. Companies' reports are obtained via a process involving recognition, adjustments and aggregation into "accounts", of all meaningful transactions

 ISBN: 978-1-61208-415-2

occurring during a given period. The resulting set of reports is made available to the public together with notes and auxiliary information.

Accounting reports are extremely efficient in revealing financial position. It is possible, for instance, to accurately predict bankruptcy more than one year before the event [14]. The direction of future earnings (up or down) is also predictable [15]. Such efficiency in conveying useful information is the ultimate reason why accounts are so often manipulated by managers.

Financial analysis of a company is typically based on the comparison of two monetary amounts (hereafter referred to as "items") taken from published reports. For instance, when a company's net income at the end of a given period is compared with assets required to generate such income, an indication of "Profitability" emerges. Pairs of items are often expressed in the form of a single value, their ratio. Since the dimension effect is similar for all items taken from the same company and period, dimension cancels out when a ratio is formed. Thus, ratios compare features such as performance of companies of different dimension [16]. Ratios are also used to detect fraud [3][4]. Indeed, most analytical tasks involving accounting information require the use of appropriately chosen ratios so that companies of different sizes can be compared while their financial features are highlighted. In this paper, an MLP training method is described whereby ratios with optimal performance characteristics are uncovered.

B. Web Mining of XBRL-encoded reports

Until recently, accounting reports were published in a variety of formats including PDF, MS Word and MS Excel. This forced users and their supporting tools into a significant amount of interpretation and manual manipulation of meta-data and led to inefficiencies and costs. From 2010 on, the Securities and Exchange Commission (SEC) of the US, as well as United Kingdom's Revenue & Customs (HMRC) and other regulatory bodies, require companies to make their financial statements public using the XML-based eXtensible Business Reporting Language (XBRL). Users of XBRL now include securities' regulators, banking regulators, business registrars, tax-filing agencies, national statistical agencies plus, of course, investors and financial analysts worldwide [17]. XML syntax and related standards such as XML Schema, XLink, XPath and Namespaces are all incorporated into XBRL, which can thus extract financial data unambiguously. Communications are defined by metadata set out in taxonomies describing definitions of reported monetary values as well as relationships between them. XBRL thus brings semantic meaning into financial reporting, promoting harmonization, interoperability and greatly facilitating ETL tasks. Web Mining of financial data is now at hand.

The initial module of the tool proposed here carries out Web Mining of XBRL content. The user first defines a selection criteria namely an industrial group, a range of assets' dimensions or simply a set of companies' codes. Then specific Web locations are searched. In the US, for instance, one such location is the SEC repository (known as EDGAR) of "fillings" of companies' reports and other data. Reports obeying stipulated criteria are downloaded and items required by the analysis are stored.

C. Data and Models

After mining and storage, three MLP are set to separately predict fraud vs no-fraud cases plus two other attributes known to widen fraud propensity space. Inputs to each of the three MLP are collections of items which were utilized as numerators or denominators of ratios in published research, namely:

- fraudulent vs non-fraudulent reports [3][4]
- bankrupted vs solvent companies [14]
- profits-up vs profits-down one year ahead [15].

Collections of items are taken from the same company reports (instance j) and may include 8 to 12 items. Both the actual period, t, and previous period, $t-1$, are collected. Items which assume positive and negative values such as net income are replaced by their two positive-only components. Input variables and target attributes used in the training and testing of the three MLP are extracted from three sources:

- "Compustat", the *de facto* repository of US companies' financial information, made available by Standard & Poor's;
- a total of 1,300 Accounting and Auditing Enforcement Releases (AAER) issued by the SEC, identifying a given set of accounts as fraudulent [4], covering the period 1983-2013, are made available by the Haas School of Business (Centre for Financial Reporting and Management) at the University of California, Berkeley;
- a list of 750 US bankruptcies covering the period 1992-2005, is made available by Professor Edward Altman from New York University.

Before training, MLP architectures consist of up to 12 inputs corresponding to collections of items just mentioned, one hidden layer with 6 nodes and two symmetrical output nodes. Hyperbolic tangents (threshold functions symmetrical around zero) are used as transfer functions in all nodes. During training, balanced matching of cases is carried out using same industry, same size (Total Assets decile) and same year companies with opposite class attribute. Training- and testing-sets are equally matched. Financial companies such as banks are excluded.

D. Knowledge Extraction

Studies on the statistical characteristics of items from accounting reports brought to light two facts. First, in cross-section the probability density function governing such items is nearly lognormal. Second, items taken from the same set of accounts share most of their variability as the dimension effect is prevalent [16]. Thus, the variability of logarithm of item i from set of accounts j, $\log x_{ij}$, is explained as the dimension effect σ_j, which is present in all items from j, plus some residual variability ε_i particular to item i:

$$\log x_{ij} = \mu_i + \sigma_j + \varepsilon_i \quad (1)$$

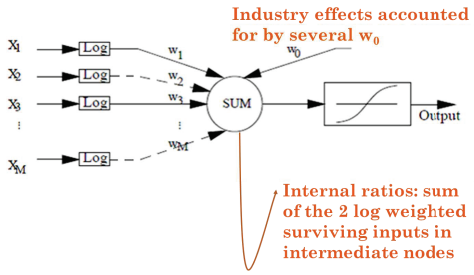

Figure 1. Ratio x_{kj}/x_{ij} of items k and i from report j is formed in MLP hidden node as a log representation $\log x_{kj} - \log x_{ij}$ because synaptic weights assume symmetrical values: $w_k = -w_i$.

μ_i is the item- and industry-specific expected value. It is thus clear why ratios formed with two items from the same set of accounts are effective in conveying financial information: the dimension effect, σ_j, cancels out when a ratio is formed. Median ratios are industry expectations while deviations from expectation observed in company j reveal how well j is doing no matter its dimension. For instance, in a given industry the median ratio of net income to assets is, say, 0.15. Any company with a ratio above 0.15, no matter small or large, is doing better than the industry.

When analyzing features such as Solvency or the likelihood of fraud, financial analysts need to know which ratios are at work, their position in relation to expectations and in which direction they are moving. In order to respond to the first of such demands, MLP training includes the competitive pruning of synaptic weights linking inputs, the $\log x_i$ in (1), to hidden nodes so that, at the end of training, only the two most relevant weights in each hidden node are permitted to survive. In a later phase, nodes also compete for survival. MLP training encompasses 5 steps:

Step	I	No penalization of synaptic weights.
	II	All hidden-node synaptic weights are equally penalized.
	III	Penalization of less relevant weights but two, one node at a time.
	IV	Zero-valued weights, all but two in each node, are pruned.
	V	Node-pruning.

In this way, internal representations similar to ratios in log space are formed inside each surviving hidden node. Here, the term "internal representation" refers to values assumed by each hidden node after summation but before transfer function, as depicted in Figure 1. The fact that each node succeeds in forming a ratio is visible through the examination of its two surviving synaptic weights: they are of similar magnitude but with opposite sign so that, after summation, a log-ratio (a difference between two $\log x_j$) is formed. Although absolute values of the two surviving weights in each hidden node are not much different from one another, they differ across nodes. Such difference reflects the importance of each node for the final classification performance.

Internal representations tend to assume the form of ratios because instances used in MLP training greatly differ in dimension while the attribute to be predicted is indeed predictable. Hidden nodes thus tend to self-organize themselves into dimension-independent variables, efficient in predicting such attribute. And since only two of the synaptic weights in each node, the most explanatory of them, are allowed to survive, weights' final values tend to assume symmetrical values so that their summation is indeed dimension-free. Representations thus mimic ratios and can be interpreted similarly.

After appropriate ratios are selected, analysts interpret their observed, company-specific deviations from industry expectation. Correspondingly, each hidden node in the MLP has a set of dummy inputs assuming the value of 1 or 0 depending on the industrial group of instance j. In this way, expected μ_i from (1) are also modelled and accounted for inside each hidden node. Since node outputs and attributes' classes are both balanced, the effect of industry dummies is to subtract industry-specific log-ratio standards from internal representations thus making them similar to a difference of two ε_i in (1). Such difference is, in log space, what analysts seek when they compare a ratio with its industry expectation.

E. Trajectories in a Features' Map

Finally, analysts observe in which direction ratios move. Internal representations are likewise input to a 2-dimensional Kohonen Features Map with 8×8 nodes. MLP outputs (transformed to become 0-1 variables) are combined with Prevalence numbers (prior probabilities of fraud) so as to approach posterior probabilities of fraud given observed features. After training, clusters are formed in the Kohonen Map, denoting identifiable features such as Solvency, Profitability, Fraud or their opposites. Visual examination of features' maps facilitates interpretation, both proximity to a given cluster and trajectories towards or away from clusters being informative.

F. Outputs to be Used by Analysts

When analysing a company's reports, analysts base their diagnostic on several concurring pieces of evidence, in favour or against *a priori* hypotheses. On the other hand, extant research on accounting manipulation suggests that fraudulent numbers lead to detectable imbalances in financial features. For instance, income may increase without the corresponding, usual increase in free cash. The selection of the two attributes complementing fraud (bankruptcy and profit direction) responds to imbalances mentioned in published research [3][4] and to the need, in the part of analysts, to examine concurring facts. Each company being analysed generates two sets of results corresponding to time periods $t-1$ and t. Output to analysts consists of the following:

1) Three posterior probabilities: fraud, default and profits going down, with a sign indicating the direction of their change from $t-1$ to t.
2) The 9 most significant values internal representations assume at period t, three from each MLP, expressed as percent increase or decrease in relation

 ISBN: 978-1-61208-415-2

to industry expectations, with a sign indicating the direction of change from $t-1$ to t.

3) Visualization of features and their trajectories from $t-1$ to t, allowing the detection of trends towards a given cluster.
4) Identification of companies neighbouring, in the features map, the company being analysed.

III. Preliminary Results, Deployment

MLP test-set performance is similar to that reported for other AI tools: 80% success in detecting fraud (6 surviving nodes), 96% success in detecting bankruptcy (5 nodes) and 78% success in predicting earnings' increase one year ahead (6 nodes). Errors are balanced: Type II error (the most expensive in this case) is reduced in relation to published research while Type I error is increased. The number of variables and synaptic weights engaged in modelling is less than half of that reported in the literature. Robustness is expected to be higher. In the downside, ratios that are formed and MLP performance both depend on broad industry type.

The tool has been set up using a variety of packages and languages; it is to be deployed as a Java-based set of modules as depicted in Figure 2. With the exception of the MLP algorithm, the analytical core will be written in R-language.

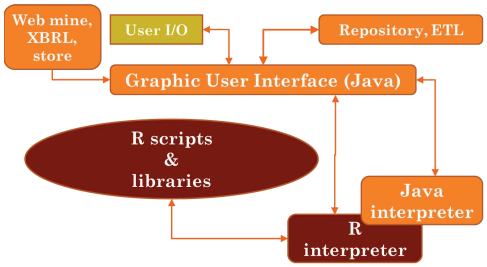

Figure 2. Architecture of the tool to be deployed.

IV. Conclusion and Future Work

Till the present date, the use of AI tools to help in the detection of manipulated accounts has been limited due to difficulties in extraction and put in place of data and also due to the "black box" nature of such tools. The present work-in-progress aims at solving both problems, producing automated and interpretable diagnostics. In the hands of analysts, the tool is self-explanatory, not just pointing out companies likely to have committed fraud but showing, rather than hiding, reasons behind such diagnostic.

The tool illustrates a case of close alignment between users' needs and algorithmic characteristics. The tool is also an example of Knowledge Extraction whereby explanatory variables are discovered amongst many candidates so that a discriminating task is carried out with optimal performance. The choice of the algorithm, the MLP, was dictated solely by its ability to form internal representations. Neither an increased performance nor the testing of novel AI techniques is the goal here. The goal is to build a usable tool, an apparently simple task but which, in this particular subject area, has eluded research effort during the last 20 years. Thus, the ultimate test is yet to be carried out, namely whether analysts will use the tool or not.

Acknowledgment

Research sponsored by the Foundation for the Development of Science and Technology of Macau, China.

References

[1] M. Nigrini, Forensic Analytics: Methods and Techniques for Forensic Accounting. John Wiley and Sons, 2011.

[2] W. Albrecht, A. C., and M. Zimbelman, Fraud Examination. Mason, OH: South-Western Cengage Learning, 2009.

[3] M. Beneish, "The Detection of Earnings Manipulation," Financial Analysts Journal, vol. 55, no. 5, 1999, pp. 24–36.

[4] P. Dechow, W. GE, C. LARSON, and R. Sloan, "Predicting Material Accounting Misstatements," Contemporary Accounting Research, vol. 28, no. 1, 2011, pp. 17–82.

[5] E. Ngai, Y. Hu, Y. Wong, Y. Chen, and X. Sun, "The Application of Data Mining Techniques in Financial Fraud Detection: a Classification Framework and an Academic Review of Literature," Decision Support Systems, vol. 50, no. 3, 2011, pp. 559 – 569.

[6] A. Sharma and P. Panigrahi, "A Review of Financial Accounting Fraud Detection Based on Data Mining Techniques," International Journal of Computer Applications, vol. 39, no. 1, 2012.

[7] E. Kirkos, S. Charalambos, and Y. Manolopoulos, "Data Mining Techniques for the Detection of Fraudulent Financial Statements," Expert Systems with Applications, vol. 32, 2007, p. 995–1003.

[8] W. Zhou and G. Kapoor, "Detecting Evolutionary Financial Statement Fraud," Decision Support Systems, vol. 50, 2011, pp. 570–575.

[9] P. Ravisankar, V. Ravi, G. Raghava Rao, and I. Bose, "Detection of Financial Statement Fraud and Feature Selection Using Data Mining Techniques," Decision Support Systems, vol. 50, no. 2, 2011, pp. 491–500.

[10] F. H. Glancy and S. B. Yadav, "A Computational Model for Financial Reporting Fraud Detection," Decision Support Systems, vol. 50, no. 3, 2011, pp. 595–601.

[11] S.-Y. Huang, R.-H. Tsaih, and F. Yu, "Topological Pattern Discovery and Feature Extraction for Fraudulent Financial Reporting," Expert Systems with Applications, vol. 41, no. 9, 2014, pp. 4360 – 4372.

[12] L. V. S.-M. K. Phua, C. and R. Gayler, "A Comprehensive Survey of Data Mining-based Fraud Detection Research," 2005, Clayton School of Information Technology, Monash University.

[13] U. Flegel, J. Vayssire, and G. Bitz, "A State of the Art Survey of Fraud Detection Technology," in Insider Threats in Cyber Security, ser. Advances in Information Security, C. W. Probst, J. Hunker, D. Gollmann, and M. Bishop, Eds. Springer US, 2010, vol. 49, pp. 73–84.

[14] E. Altman, "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy," The Journal of Finance, vol. 23, no. 4, Sep. 1968, pp. 589–609.

[15] J. Ou and S. Penman, "Financial Statement Analysis and the Prediction of Stock Returns," Journal of Accounting and Economics, vol. 11, no. 4, 1989, pp. 295–329.

[16] S. McLeay and D. Trigueiros, "Proportionate Growth and the Theoretical Foundations of Financial Ratios," Abacus, vol. XXXVIII, no. 3, 2002, pp. 297–316.

[17] T. Dunne, C. Helliar, A. Lymer, and R. Mousa, "Stakeholder Engagement in Internet Financial Reporting: The Diffusion of {XBRL} in the {UK}," The British Accounting Review, vol. 45, no. 3, 2013, pp. 167 – 182.

 ISBN: 978-1-61208-415-2