A Mixed-Methods Approach to Poor Disclosure Detection

Eric R. Pearson[1] **Eden Smith**[2] **Jinming Zhao**[1] **Chan Cheah**[1] **Wray Buntine**[1] **Ann E. Nicholson**[1]

[1]Faculty of Information Technology, Monash University, Clayton, Victoria, AU
[2]School of Historical and Philosophical Studies, University of Melbourne, Melbourne, Victoria, AU

1 EXTENDED ABSTRACT

1.1 INTRODUCTION

The Australian Securities and Investment Commission (ASIC) is responsible for ensuring the integrity, fairness and efficiency of Australia's financial markets, which is underpinned by investors having the confidence in accessing transparent and equitable market information that informs their decision making. While responsible oversight provides strict disclosure rules, securities commission's typically lack the resources required to discover withheld, misstated or mistimed adverse information in listed companies' market announcements and public disclosures [Ana Carvajal and Jennifer Elliott, 2007]. Providing a digital capacity for regulators to detect and quickly respond to listed companies disclosure misinformation or disinformation incidents has the potential to offset these resource limitations. Advances in technology including Artificial Intelligence (AI) and machine learning oriented approaches to unstructured data processing have the potential to significantly improve a security commission's capacity to detect poor disclosure behaviour. During a 3-month project we developed a novel mixed-methods approach to both automate disclosure classification as well as identify poor disclosure behaviours. These methods included qualitative analysis (QA), natural language programming (NLP) and Bayesian Networks (BN).

1.2 APPROACHES

Qualitative Analysis (QA) is the analysis of qualitative data, for example unstructured human text. Content analysis is one of the more widely used qualitative methods for studying the disclosure statements of companies [Khandelwal et al., 2021]. This approach involves identifying the appearance of content-types within text using a set of analytic categories ('codes') [Hsieh and Shannon, 2005]. Those analytic-categories can be formalised by developing a set of guidelines ('codebook') for consistent annotation. A codebook helps to increase the rigor and reproducibility of this type of qualitative approach [Roberts et al., 2019, Popping, 2000]. Natural Language Processing (NLP) refers to technology that allows computers to analyse large volumes of unstructured text data [Jurafsky and Manning, 2012]. NLP can help businesses automate processes and reduce human errors [Liu et al., 2021] as well as support decision making [Chen et al., 2020]. We used two well known NLP toolkits during this study, Google BERT [Devlin et al., 2018] and IBM Watson[1]. Since its introduction in 2018, BERT and its variants have achieved great success in various classification tasks [Radford et al., 2019, Liu et al., 2021, Brown et al., 2020]. While vanilla BERT processes paragraphs and sentences, it can be extended to processing long documents for classification purposes [Adhikari et al., 2019]. IBM Watson is a well-established business intelligence (BI) product that provides a strong analytics engine together with a natural language querying tool. For example, it is empowered to analyse text to extract metadata from content (e.g., emotion, sentiments). The toolkit has been extensively used in the industry for its accessible user interface and robust language querying functions [Ferrucci, 2012]. Bayesian Networks (BNs) are directed acyclic graphs (DAG) that are underpinned by conditional probability distributions [Korb and Nicholson, 2010]. Bayesian Networks offer advantages over contemporary ML techniques in that they can be created where there may be little or no existing data (only expert knowledge) [de Waal et al., 2016], where the problem is inherently uncertain [Iqbal et al., 2015], and in the specific case of anomaly detection, there are few anomalous events to draw upon [Mascaro et al., 2010, 2014, Hauskrecht et al., 2007, Hill et al., 2009]. BNs are particularly suited to the task of anomaly detection as they don't necessarily need prior knowledge of an anomaly to function [Zhang et al., 2009].

[1]https://www.ibm.com/au-en/watson

Accepted for the 16th Bayesian Modelling Applications Workshop (@UAI2022) (BMAW 2022).

1.3 METHODS

We conducted systematic content-analysis to develop a set of analytic categories for identifying characteristics of financial disclosure statements that, when integrated within NLP and BN models, contribute to identifying anomalous company behaviours. In total 26 characteristics were identified. Instructions for systematically identifying these characteristics were documented within a Codebook that was used to annotate 243 disclosure statements. NLP was applied to automate the classification of these characteristics. Individual disclosures range from less than a page to more than a hundred pages in length and as such required a long text classification solution. Google BERT was selected for its demonstrated long text classification capabilities [Devlin et al., 2018]. Qualitative attributes were augmented with disclosure meta-characteristics such as sentiment and emotion using IBM Watson. We limited our sample set to only 30,000 of the potential millions of disclosures (disclosures with less than 5 pages and belonging to the same disclosure category). Naive bayes BN models were constructed using both the qualitative characteristics and meta characteristics. The naive bayes models were trialed with 3 different classifiers all of which were proxies for companies that conducted poor disclosure behaviour. We conducted a sensitivity analysis on the networks to understand which characteristics were most influential towards the classifier. Our final modelling step was to group companies and their disclosures into cohorts and identify characteristics that were at least 2 standard deviations from the norm as anomalies. A future step will be to conduct a correlation study on these anomalies to determine if a relationship exists between anomalous characteristics and a company's poor disclosure behaviour.

Figure 1 details the mixed methods process flow. Activities in blue are data wrangling, green were conducted by our qualitative team, orange were our NLP processes and purple our BN and main output activities.

A brief explanation of this diagram is provided. The process starts with the literature review having information sourced using Google Scholar, and the ASX and ASIC websites. ASX provided 3 years of disclosure history that we imported into an AWS database. The information from the literature review, raw data and knowledge from industry experts allowed us to define the data in a data dictionary as well as build the Codebook for the qualitative processes. From these sources we also built our NLP and BN models. We then qualitatively coded a sample of 243 disclosures that were then used as the bases for training our NLP qualitative classification model. Both the human labelled and NLP labelled disclosures were then used to train the BN models. Finally we conducted validation as well as a case study review that also considered AFR, Twitter and ASX market activity. Each model, and the data wrangling and cleansing tasks were designed with automation in mind. This allowed

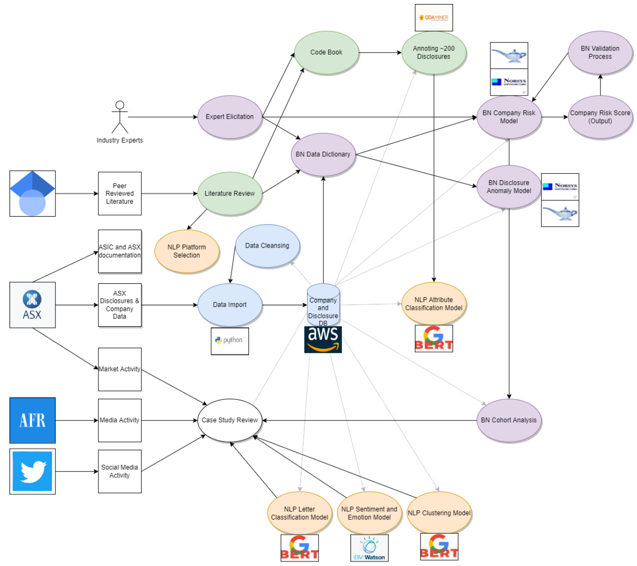

Figure 1: Mixed methods process flow.

for model rapid iteration and trialling of different combinations of model type, classifier and source data. A second rationale for the rapid interaction was the agile nature of the feasibility study and the fact that modelling and pipeline development were happening in parallel to knowledge engineering. This resulted in model design and data schema that changed frequently.

1.4 CONCLUSION

The three approaches of QA, NLP and BN combined in a mixed-methods approach were demonstrated to streamline both human led qualitative processes and AI automation. The iterative approach allowed for rapid change management, data changes and model tuning. Furthermore the BN models ranked companies' risk of poor disclosure as well as identified anomalies in company disclosure behaviour. Further work needs to be undertaken, particularly in coding more disclosures and exploring the relationships between anomalous behaviours and poor disclosure behaviour.

References

Ashutosh Adhikari, Achyudh Ram, Raphael Tang, and Jimmy Lin. Docbert: Bert for document classification. *arXiv preprint arXiv:1904.08398*, 2019.

Ana Carvajal and Jennifer Elliott. Strengths and Weaknesses in Securities Market Regulation: A Global Analysis, November 2007. URL `https://www.imf.org/external/pubs/ft/wp/2007/wp07259.pdf`.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and others. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. NLP in FinTech applications: past, present and future. *arXiv preprint arXiv:2005.01320*, 2020.

Alta de Waal, Hildegarde Koen, Pieter de Villiers, Henk Roodt, Nyalleng Moorosi, and Gregor Pavlin. Construction and evaluation of Bayesian networks with expert-defined latent variables. In *2016 19th International Conference on Information Fusion (FUSION)*, pages 774–781. IEEE, 2016.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

David A Ferrucci. Introduction to “this is watson”. *IBM Journal of Research and Development*, 56(3.4):1–1, 2012. Publisher: IBM.

Milos Hauskrecht, Michal Valko, Branislav Kveton, Shyam Visweswaran, and Gregory F. Cooper. Evidence-based anomaly detection in clinical domains. In *AMIA Annual Symposium Proceedings*, volume 2007, page 319. American Medical Informatics Association, 2007.

David J. Hill, Barbara S. Minsker, and Eyal Amir. Real-time Bayesian anomaly detection in streaming environmental data. *Water Resources Research*, 45(4), 2009. Publisher: Wiley Online Library.

Hsiu-Fang Hsieh and Sarah E. Shannon. Three approaches to qualitative content analysis. *Qualitative health research*, 15(9):1277–1288, 2005. Publisher: Sage Publications Sage CA: Thousand Oaks, CA.

Khalid Iqbal, Xu-Cheng Yin, Hong-Wei Hao, Qazi Mudassar Ilyas, and Hazrat Ali. An overview of bayesian network applications in uncertain domains. *International Journal of Computer Theory and Engineering*, 7(6):416, 2015. Publisher: IACSIT Press.

Dan Jurafsky and Christopher Manning. Natural language processing. *Instructor*, 212(998):3482, 2012.

Chandani Khandelwal, Satish Kumar, and Deepak Verma. Exploring financial risks disclosure: evidence from Indian listed companies. *Qualitative Research in Financial Markets*, 2021. Publisher: Emerald Publishing Limited.

Kevin B. Korb and Ann E. Nicholson. *Bayesian Artificial Intelligence*. CRC Press, December 2010. ISBN 978-1-4398-1592-2. Google-Books-ID: LxXOBQAAQBAJ.

Zhuang Liu, Degen Huang, Kaiyu Huang, Zhuang Li, and Jun Zhao. Finbert: A pre-trained financial language representation model for financial text mining. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 4513–4519, 2021.

Steven Mascaro, Kevin B. Korb, and Ann E. Nicholson. Learning abnormal vessel behaviour from ais data with bayesian networks at two time scales. *Tracks a Journal of Artists Writings*, pages 1–34, 2010.

Steven Mascaro, Ann E. Nicholso, and Kevin B. Korb. Anomaly detection in vessel tracks using Bayesian networks. *International Journal of Approximate Reasoning*, 55(1):84–98, 2014. Publisher: Elsevier.

Roel Popping. Network text analysis. *Computer-Assisted Text Analysis, Sage, London*, pages 97–128, 2000.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and others. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

Kate Roberts, Anthony Dowell, and Jing-Bao Nie. Attempting rigour and replicability in thematic analysis of qualitative research data; a case study of codebook development. *BMC medical research methodology*, 19(1):1–8, 2019. Publisher: Springer.

Weiyu Zhang, Qingbo Yang, and Yushui Geng. A survey of anomaly detection methods in networks. In *2009 International Symposium on Computer Network and Multimedia Technology*, pages 1–3. IEEE, 2009.