Sharpening Ponzi Schemes Detection on Ethereum with Machine Learning

Oral Communication

Letterio Galletta Fabio Pinelli
IMT School for Advanced Studies Lucca, Lucca, Italy

***Abstract*—Blockchain technology has been successfully exploited for deploying new economic applications. However, it has started arousing the interest of malicious actors who deliver scams to deceive honest users and to gain economic advantages. Ponzi schemes are one of the most common scams. Here, we present a work in progress to build a classifier for detecting smart Ponzi contracts on Ethereum, which can be used as the backbone for developing detection tools.**

***Index Terms*—blockchain fraud, Ponzi scheme, smart contracts**

I. Introduction

Blockchain is revolutionizing how individuals and companies exchange digital assets without the control of a central authority. This technology has been successfully exploited for deploying new economic applications, e.g., cryptocurrencies [1] and DeFi [2]. However, soon after this technology became widespread and its economic value increased, it has started arousing the interest of malicious actors who are eager to take some advantages due to the pseudonymity of these platforms and the lack of regulation [3]: on the one hand, they exploit cryptocurrencies to transfer currency without being tracked by authorities; on the other hand, they deliver scams to deceive honest users willing to make revenues through cryptocurrencies. Nowadays, many types of scams can be found on blockchain platforms, such as exploits, hacks, and phishing [4]: estimates say that scams in Bitcoin [5] gathered more than 7 million USD.

Among the various scams, Ponzi schemes have approached the blockchain world, first on Bitcoin [5] and more recently on Ethereum [6]. These are fraudulent investment operations where older investors obtain returns from new investors' money rather than legitimate business activities. Although the actual conditions to gain money depend on the specific rules of the scheme, a common feature is that participants who want to redeem their investments have to make new participants join the scheme. Participants who join later are the most likely to lose their money. Thus, the development of automatic techniques to counter these scams is required to protect average users and to allow them to participate safely in the blockchain economy.

This paper briefly describes a work in progress to build an automatic technique for classifying smart contracts, which can be used as the backbone for developing new detection tools. Here, we focus on Ethereum and smart contracts to deliver Ponzi schemes, called smart Ponzi contracts. Since the entire transaction history and the contracts' bytecodes are publicly available and provide accurate records of user and contract behaviours, machine learning techniques are a natural choice to detect possible frauds and scams. More precisely, we aim at a threefold contribution. First, we are addressing the problem of the unavailability of public data sets to train effective automatic classifiers. Even if blockchain data is publicly available, the literature is missing a reference dataset of Smart Ponzi contracts. Many papers in the literature train classifiers on their own datasets with their own features, but they are not publicly available. This makes it difficult to compare different proposals because the performances of a classifier depend on the data used to train it. For this reason, we release a reusable data set that collects 4422 unique real-world smart contracts, where 3749 (84.78%) are not-Ponzi, and 673 (15.22%) are Ponzi. Our data set contains both information about the transaction history of the contracts as well as their bytecode. Another issue we are investigating is defining a small and effective set of features that ensures an accurate binary classification process. To find this set, we proceed as follows. First, we consider the four requirements proposed by Bartoletti et al. [6] to classify a smart contract as a Ponzi scheme and consider sets of features proposed in the literature and check that they do not capture well some aspects of such requirements. Thus, we introduce new features that fill this gap and show how they improve the classification through experiments. Since we want to maintain the set of features as small as possible, also we identify those that can be removed since their contribution to the classification is minimal. We adopt eXplainable AI (XAI) techniques to investigate the contribution of each feature. Finally, as result of this process, we obtained an accurate binary classification model to detect smart Ponzi contracts. Moreover, our experiments show that the proposed model performs better than the ones proposed in the literature when considering the AUC as a metric and achieves high accuracy for practical use.

In summary, our contributions are:

- a reusable and publicly available data set of 4422 real-world smart contracts where 3749 are not Ponzi, and 673 are Ponzi;
- a small and effective set of features that ensures a good classification quality;

- a binary classifier to detect smart Ponzi contracts that outperform classifiers in the literature when considering the AUC as a metric.

We proceed as follows. Section II briefly describes our dataset and sketches the methodology used to build our binary classifier and to study its quality and the impact and importance of the features. Section III concludes the paper by discussing future work. The data set and the notebooks used for the experiments presented in this paper are available online[1]. A full version of this paper is available on arxiv [7].

II. Dataset and experiments

We built our data set based on others from the literature [6], [8], [9], extending the set of features trying to satisfy the requirements by Bartoletti et al. [6] and updating the blockchain data. As for the previous papers, we manually label the dataset by inspecting the contract code and check if it satisfies Bartoletti et al.'s requirements. The resulting dataset contains 4422 smart contracts, with 3749 (85.23%) labelled as not-Ponzi and 673 (14.77%) as Ponzi. Actually, from this source dataset, we build three datasets D1 D2 and D3: D1 uses all the features, D2 uses only features from the literature, whereas D3 contains only the features that provide the best classification in our experiments.

To study if the new features improve the classification, we select the best classifiers on D1 and D2 and compare their performances. In particular, we consider Decision Tree [10], Random Forest [11], and Light Gradient Boosting Machine Classifier (LGBMC) [12] as classifiers and perform a grid search procedure with cross-validation to fine-tune the hyper-parameters of each classifier optimizing the AUC metric. Once we have selected the best values for the hyper-parameters, we compute the standard metrics *Accuracy*, *AUC*, *F1*, *Precision*, and *Recall* on the test set for each classifier. According to the AUC metric, the best model for both datasets is LGBMC. Then, we compute the diagnostic abilities of the resulting classifiers by studying their ROC curve that describes the true positive and the false positive rate. From the ROC curve, the classifier trained on D1 performs better than the one trained on D2.

Once we select the best classifier, we investigate the contribution to the classification for each feature. We proceed as follows. First, we take our best model on D1 and determine the importance of each feature. Then, we determine if there exists a subset of the features of D1 that improves the quality of the classifier. To achieve that, we consider the number of features like another hyper-parameter and perform a grid search procedure with cross-validation to optimise it with the AUC metric. Our tuning procedure works as follows. We start considering all the features in D1, perform a grid search and 5-fold cross-validation, and aim to optimise the AUC. The result of this step is the best-performing classifier on the current set of features. Then, we adopt the *Recursive Feature Elimination* algorithm to remove the less important feature. Given the new reduced data set, the procedure is repeated until we obtain a classifier with a worse value of AUC. In our experiment, the data set with the highest mean AUC is D3 which includes only 25 features. Hence, the LGBMC trained on D3 performs better and improves the best classifier on D1.

[1] https://github.com/fpinell/ponzi_ml

III. Conclusion

This paper presented a brief overview of a work in progress towards an automatic technique for detecting smart Ponzi contracts on Ethereum. We released a reusable data set with 4422 unique real-world smart contracts. Then, we introduced a new set of features that allowed us to improve the classification. In the full version of this paper [7], we show that our classifier outperforms previous efforts in the literature [8]. Finally, we experimentally identified a small and effective set of features that ensures a good classification quality.

In future work, we plan to extend our investigation towards different directions. First, we intend to improve the procedure to optimize the best set of features. Then, we would also like to consider the bytecode of contracts present in our dataset but not currently used by our classifier. Our idea is to derive some code features that allow us to reduce the blockchain features and improve the classification performances. Moreover, we plan to apply deep learning techniques to minimize the feature engineering effort, especially in the presence of bytecode. Finally, we plan to study whether our approach can be applied to detect other forms of scams on Ethereum, and phishing is one of the most promising.

References

[1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," https://bitcoin.org/bitcoin.pdf, 2008.

[2] E. Napoletano, "Decentralized finance is building a new financial system," https://www.nasdaq.com/articles/decentralized-finance-is-building-a-new-financial-system-2021-04-02, 2021, (last access 2022).

[3] T. Moore, "The promise and perils of digital currencies," *International Journal of Critical Infrastructure Protection*, vol. 6, no. 3, pp. 147–149, 2013.

[4] M. Bartoletti, S. Lande, A. Loddo, L. Pompianu, and S. Serusi, "Cryptocurrency scams: Analysis and perspectives," *IEEE Access*, vol. 9, pp. 148 353–148 373, 2021.

[5] M. Vasek and T. Moore, "There's no free lunch, even using bitcoin: Tracking the popularity and profits of virtual currency scams," in *Financial Cryptography and Data Security*. Springer, 2015, pp. 44–61.

[6] M. Bartoletti, S. Carta, T. Cimoli, and R. Saia, "Dissecting ponzi schemes on ethereum: Identification, analysis, and impact," *Future Gener. Comput. Syst.*, vol. 102, pp. 259–277, 2020.

[7] L. Galletta and F. Pinelli, "Sharpening ponzi schemes detection on ethereum with machine learning," 2023.

[8] W. Chen, Z. Zheng, E. C.-H. Ngai, P. Zheng, and Y. Zhou, "Exploiting blockchain data to detect smart ponzi schemes on ethereum," *IEEE Access*, vol. 7, pp. 37 575–37 586, 2019.

[9] W. Chen, X. Li, Y. Sui, N. He, H. Wang, L. Wu, and X. Luo, "Sadponzi: Detecting and characterizing ponzi schemes in ethereum smart contracts," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 5, no. 2, 2021.

[10] J. R. Quinlan, "Induction of decision trees," *Machine learning*, vol. 1, no. 1, pp. 81–106, 1986.

[11] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[12] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "Lightgbm: A highly efficient gradient boosting decision tree," *Advances in neural information processing systems*, vol. 30, 2017.