Prerequisites and AI challenges for model-based Anti-Money Laundering

Marvin Oeben[1] , **Jeroen Goudsmit**[2] and **Elena Marchiori**[1]
[1]Radboud Universiteit Nijmegen
[2]Vrije Universiteit Amsterdam
{m.oeben, e.marchiori}@cs.ru.nl, j.p.goudsmit@vu.nl

Abstract

Money laundering is a large societal problem. Anti-money laundering is arguably ineffective and knows many challenges. In this position paper we highlight prerequisites for comparable model-based anti-money laundering, indicate whether these are met, and make recommendations on how to further this field in both a fundamental as well as an experimental manner.

1 Money Laundering as a societal problem

Money laundering allows criminals to introduce the proceeds of crime into the legitimate economy. It is one of the key engines of organised crime and is linked to virtually all criminal activities generating proceeds [European Police Office, 2017]. Reducing these activities is included as the fourth target under the 16th Sustainable Development Goal: *Peace, Justice and Strong Institutions*. But how do you detect criminal proceeds disguised as legitimate money? This paper focuses on model-based efforts for anti-money laundering (AML) and provides recommendations for research in this area.

The current estimate for the amount of laundered money is between 2% and 5% of the global GDP, this translates to between 1.7 and 4.5 trillion USD today [International Monetary Fund, 2019], of which less than 1% is seized and frozen [United Nations Office on Drugs and Crime, 2011]. This small number of seizures means that we have little information on the total problem of money laundering [Levi *et al.*, 2018].

It is not straightforward to define or categorise money laundering activities. Many criminals are highly flexible and display great adaptability and speed in which they adjust their *modus operandi* [European Police Office, 2017]. Money laundering is generally viewed as a three-phase process: placement, layering, and integration [Financial Action Task Force, 2018; Cox, 2014]. In placement, illegally obtained money enters the financial system. Through layering, its sources are consequently obfuscated, hiding the illegitimacy of its origins. Finally, in integration, the funds are integrated into the legitimate economy — for example as investments in real estate or luxury goods.

The United Nations has criminalised money laundering and has given states the obligation to install measures to combat money laundering [United Nations, 2000]. Agents in the financial system typically have a legal obligation to carry out monitoring of transactions to enable the detection of suspicious transactions. For example, the EU Anti-Money Laundering Directive provides both objective and subjective criteria under which an *obliged entity* (including but not limited to financial institutions, accountants, and notaries) is under a legal obligation to report a financial transaction to the competent authorities [European Parliament, 2018]. Objective criteria are crisp and clear: for example, cash deposits exceeding $10K in value. Subjective criteria are open to interpretation and rely on professional judgement following a risk-based approach [Joint Committee, 2018].

In this paper, we discuss the model-based monitoring of transactions and the subsequent detection of money laundering. In section 2, we argue for three key prerequisites for model-based AML: (2.1) agree on a method of fairly comparing models, (2.2) agree on validation of models and results and (2.3) agree on data points that represent group behaviour over time. We stress that these points are model-agnostic: they are the prerequisites for mature model development. It is our position that none of these prerequisites are currently being met. In section 3, we provide recommendations on (3.1) how financial institutions and states through experimentation and (3.2) academics through fundamental research can improve this situation even when the prerequisites are not yet being met.

2 Prerequisites for model-based AML

Model-based AML amounts to a classification problem: determine whether a transaction meets the legally defined criteria for reporting suspicious behaviour. Such models have been developed, focusing on the perspective of an individual *obliged entity* [Chen *et al.*, 2017; Salehi *et al.*, 2017]. Below we discuss three criteria that we consider prerequisites for the development of more mature model-based AML. It is our position that these criteria are not met at the time of writing.

2.1 Agree on a method of fairly comparing models

Without a fair comparison of model performance the academic community can not assess the merits of a new model compared to that of prior models. Such a comparison in essence amounts to evaluating similar models along similar data sets. This can be achieved in two fundamentally distinct ways: either by exchanging models or exchanging data sets.

Exchanging data sets is likely to be challenging. Indeed, the sensitive nature of the data involved prohibits it from being readily shared among researchers in general. A benchmark data set of true transactions, however useful, would be difficult to compose without either compromising the privacy of the entities involved or damaging usability through anonymisation. In 3.2 we discuss the option of generating synthetic data to be used as a benchmark across the industry.

We recommend building a framework to exchange AML models for the purposes of joint comparison. Such a framework ought to be accessible to obliged entities who may use this framework to submit their model for evaluation across all other entities. By running the model across the data sets at all entities in the framework, one can evaluate its outcome against all other prior submitted models. This allows for a collective comparison, facilitating further growth. Even if this framework is not available, obliged entities can start comparing models and experimenting with possibilities as noted in section 3.1.

2.2 Agree on validation of models and results

Model-based AML is designed to fulfill the legal obligation to detect and report suspicious activity by obliged entities. The reported activities are by no means iron clad cases of money laundering leading to a criminal conviction. Indeed, not all behaviour reported by obliged entities is brought to trial and not all such trials lead to a criminal conviction [ACAMS and Dow Jones, 2016]. Moreover, money laundering may very well not be detected nor reported as suspicious activity.

In order to evaluate the performance of an AML model, one ought to assess whether any given transaction could be construed as suspicious. Reasonable people may disagree on the subjective interpretation of this definition. Indeed, two organisations may use two different standards of "suspicious" and thus work with two very different models. It is difficult to interpret any comparison between these models, as neither definition of "suspicious" is *a priori* better than the other. It may very well be that one definition is stricter but fails to identify any additional cases of money laundering. Moreover, it is meaningless to say that one model fares better than another, simply because it classifies more transactions as suspicious. A golden standard for comparing notions of suspicious is needed to resolve this unclarity.

Any suspicious activity detected by an obliged entity is to be reported to its Financial Intelligence Unit [United Nations, 2000]. If this entity were to receive a report on which of these transactions actually constitute money laundering, then it could use this information to validate its notion of suspicious. We recommend that such a report be given for a significant sample of reported suspicious transactions. This would allow the validation of models and their results.

Alternatively, objective criteria can be composed to approximate the current subjective criteria as discussed in 3.1. This would amount to defining known typologies in mathematical terms, with the obvious limitation that unknown typologies will be missed.

2.3 Agree on data-points that represent behaviour of networks over time

Conventional AML systems are implemented as expert systems which employ risk-based scenarios and thresholds [Unger and Van Waarden, 2009]. These systems are a good way to model the objective reporting criteria, as regulation can be directly translated into knowledge for the system. For the subjective criteria however, they suffer from a large number of false positives [ACAMS and Dow Jones, 2016]. When trying to improve on this, agreement needs to be reached on which data have to be available and which information should be used, in order to create a proper model-based solution.

Money laundering spans networks

Most of the money laundering is performed by organised crime groups (OCGs), of which the majority operates in a hierarchical setting [European Police Office, 2017]. An OCG is a group of three or more persons existing over a period of time acting in concert with the aim of committing crimes for financial or material benefit [Council of the European Union, 2008]. As these OCGs operate in criminal networks, it is preferable to be able to identify not only actors or transactions but also networks of malicious intent. Yet, only one third of the surveyed literature has considered network enhanced models [Chen *et al.*, 2017]. We recommend that AML models include network-based features, such as in [Colladon and Remondi, 2017; Savage *et al.*, 2016].

Money laundering takes time

Money launderers are rarely caught due to a single transaction [European Police Office, 2017]. The activity of laundering takes time, progressing through the aforementioned three stages. Conventional systems consider aggregation over time in their detection scenarios to make use of this observation [Unger and Van Waarden, 2009]. However, one third of the surveyed machine learning literature on AML does not have a time component [Chen *et al.*, 2017]. We thus recommend that AML models consider transactions over time as opposed to singular transactions.

3 Directions for further AML research

When looking at the main prerequisites mentioned in the previous section, we note that industry and regulators play a crucial role. Agreements between industry and regulators on the one side and AI researchers interested in tackling the AML problem on the other side should be stipulated. This would allow for the generation of well defined benchmark problems with corresponding publicly available data sets. Besides this ambitious plan, there are more pragmatic possibilities for AI in AML, which we summarise below. In this we mention fundamental problems for both AI researchers and industry and regulators, and explore more experimental research for researchers with access to financial data.

3.1 Experimental research on comparing and validating AML models

Even when the prerequisites in 2.1 and 2.2 are not yet fulfilled, some experimental work can be done. This work will have to be combined with agreements on the approach an FIU takes in evaluating money launder behaviour.

The first direction is in the work of OCG subgraph definition, generation and artificial addition to financial systems as closed-world examples. Once money laundering definition and (some) typologies have been agreed upon, known behaviour can be transformed into subgraphs showing said behaviour. Experimenting with artificial injection into financial systems can teach us more about situations where current detection engines may fail.

The second direction lies in experimentation with graph generation frameworks inside obliged entities. Using graph generation frameworks, a benchmark data set can be used which may help AI researchers create and enhance current classification approaches inside obliged entities and law enforcement. Research is however needed from the previous direction (for injection of OCGs) and in building a governance framework which ensures that the graph is sufficiently different from the original graph to safeguard privacy of the obliged entities' customers.

3.2 Fundamental research on network-based modelling

The basis of any new solution is fundamental research, we highlight topics that are either required (as highlighted in the previous section) or show great promise in solving some of the fundamental problems.

Graph Neural Networks

In the last decade, deep artificial neural networks have proven to be superior in many pattern classification tasks [Schmidhuber, 2015]. In recent years, graph neural networks have been an effective way of learning representations of structured data for graph classification tasks [Wu *et al.*, 2019]. As we have highlighted in 2.3 and the fact that the largest portion of Money Laundering is done by OGCs, classifiers based on network features show the most promise. There are however several interesting research questions which still need to be answered before application in AML become possible.

The first challenge lies in scalability. Even though more scalable graph neural networks have emerged recently, all are missing a property needed for successful AML on the scale of a financial system. They only provide node embeddings [Lerer *et al.*, 2019], do not incorporate directed graphs [Hamilton *et al.*, 2017] or can not yet incorporate edge information (such as transaction amounts) [Zhang *et al.*, 2018].

Current graph neural network implementations only label either nodes or complete graphs where many AML systems focus on labeling transactions (edges in a graph). Therefore, there is a challenge in exploring edge-labelling that would help the AML effort. Given the previous section, it would be even better to create a subgraph classification algorithm which can be trained to detect OCGs directly.

Graph Generation

With the increased research interest of graph neural networks, graph generators have also been proposed [Wu *et al.*, 2019]. They could be of great value in AML as they can be used to create synthetic data which is currently unavailable. In the following section we will discuss two examples (which in our eyes show the most promise for future application in AML) and highlight fundamental open problems that require further research before they can be used for our purposes.

The first graph generator is NetGAN [Bojchevski *et al.*, 2018]. This generator combines random walks, LSTM and a Wasserstein GAN [Arjovsky *et al.*, 2017] in order to generate synthetic random walks in a new graph which become indistinguishable from random walks in the original graph. This creates a generative model of (artificial) graphs encompassing properties of the original real-life graph. A big advantage of NetGAN is that it can generate a synthetic graph from just one sample. However, it still requires more research to be a good generator for AML data. It currently works on undirected graphs, where the direction of a transaction is essential information in a financial system. Even though it is trivially parallelisable, it will require extensions to make it fit for multi-million nodes and multi-billion edge graphs. It can currently only generate plain graphs, extension for attributes and addition and deletion of nodes/edges over time is required for AML research.

The second generator we highlight is Graphite [Grover *et al.*, 2018]. This model combines graph neural networks with graph kernels [Shervashidze *et al.*, 2011] and variational autoencoders [Kingma and Welling, 2013]. In contrast to NetGAN, Graphite is trained on multiple graphs, but can also work by considering sampled connected subgraphs to generate a synthetic graph from a single other graph. Moreover, it can handle attributed and directed graphs. By leveraging Monte Carlo methods, it is easily scalable to graphs of thousands of nodes. Larger size graphs will still need further research. Also, no time component is yet known for these graphs which handle node and edge addition/deletions.

The last missing component for all graph generators is how to deal with imbalancedness in the label distribution of the graph. In reality, OCGs and money launderers will only make up a small portion of the graph. Therefore, in order to generate a full financial system, a way of artificially adding subgraphs with certain behaviour will be required. To our knowledge, no research on this has been done for these models.

4 Conclusion

Money laundering is a large-scale societal problem, with total money laundered estimated to be between 2% and 5% of global GDP. Yet, it is not straightforward to identify suspicious transactions because of various reasons ranging from inability to share data to the subjective interpretation of the term "suspicious" itself. In this paper we have identified three prerequisites which need to be met before proper model-based AML research can be executed, namely, agreement on methods of fairly comparing models, on validation of the models and results and agreement on which data-points represent behaviour of networks over time. Lastly we have shown where core fundamental and experimental challenges lie for industry, regulators and AI researchers and practitioners both inside and outside of industry. We have highlighted that due to the recent discoveries in graph neural networks, artificial data generation and graph classification is in our opinion the most promising approach in the domain of AML.

References

[ACAMS and Dow Jones, 2016] ACAMS and Dow Jones. Global anti-money laundering survey results 2016, 2016.

[Arjovsky *et al.*, 2017] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.

[Bojchevski *et al.*, 2018] Aleksandar Bojchevski, Oleksandr Shchur, Daniel Zügner, and Stephan Günnemann. Netgan: Generating graphs via random walks. *arXiv preprint arXiv:1803.00816*, 2018.

[Chen *et al.*, 2017] Zhiyuan Chen, Ee Na Teoh, Amril Nazir, Ettikan Kandasamy Karuppiah, Kim Sim Lam, et al. Machine learning techniques for anti-money laundering (AML) solutions in suspicious transaction detection: a review. *Knowledge and Information Systems*, 57:1–41, 2017.

[Colladon and Remondi, 2017] Andrea Fronzetti Colladon and Elisa Remondi. Using social network analysis to prevent money laundering. *Expert Systems with Applications*, 67:49–58, 2017.

[Council of the European Union, 2008] Council of the European Union. Council Framework Decision 2008/841/JHA of 24 October 2008 on the fight against organised crime. *Official Journal of the European Union*, L300:42–45, 2008.

[Cox, 2014] D.W. Cox. *Handbook of Anti-money Laundering*. Wiley, 2014.

[European Parliament, 2018] European Parliament. Directive (EU) 2018/843 of the European Parliament and of the Council of 30 May 2018 amending Directive (EU) 2015/849 on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing, and amending Directives 2009/138/EC and 2013/36/EU. *Official Journal of the European Union*, L156:43–74, 2018.

[European Police Office, 2017] European Police Office. *Europol Serious and Organised Crime Threat Assessment*. 2017.

[Financial Action Task Force, 2018] Financial Action Task Force. *The FATF Recommendations: International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation*. oct 2018.

[Grover *et al.*, 2018] Aditya Grover, Aaron Zweig, and Stefano Ermon. Graphite: Iterative generative modeling of graphs. *arXiv preprint arXiv:1803.10459*, 2018.

[Hamilton *et al.*, 2017] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 1024–1034. Curran Associates, Inc., 2017.

[International Monetary Fund, 2019] International Monetary Fund. *World Economic Outlook: Growth Slowdown, Precarious Recovery*. Washington, DC, April 2019.

[Joint Committee, 2018] Joint Committee. Joint Guidelines on Risk Factors, Jan 2018. Joint Guidelines under Articles 17 and 18(4) of Directive (EU) 2015/849 on simplified and enhanced customer due diligence and the factors credit and financial institutions should consider when assessing the money laundering and terrorist financing risk associated with individual business relationships and occasional transactions.

[Kingma and Welling, 2013] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[Lerer *et al.*, 2019] Adam Lerer, Ledell Wu, Jiajun Shen, Timothée Lacroix, Luca Wehrstedt, Abhijit Bose, and Alexander Peysakhovich. PyTorch-BigGraph: A Large-scale Graph Embedding System. *CoRR*, abs/1903.12287, 2019.

[Levi *et al.*, 2018] Michael Levi, Peter Reuter, and Terence Halliday. Can the AML system be evaluated without better data? *Crime, Law and Social Change*, 69(2):307–328, Mar 2018.

[Salehi *et al.*, 2017] Ahmad Salehi, Mehdi Ghazanfari, and Mohammed Fathian. Data Mining Techniques for Anti Money Laundering. *International Journal of Applied Engineering Research*, 12(20):10084–10094, 2017.

[Savage *et al.*, 2016] David Savage, Qingmai Wang, Pauline Chou, Xiuzhen Zhang, and Xinghuo Yu. Detection of money laundering groups using supervised learning in networks. *arXiv preprint arXiv:1608.00708*, 2016.

[Schmidhuber, 2015] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

[Shervashidze *et al.*, 2011] Nino Shervashidze, Pascal Schweitzer, Erik Jan van Leeuwen, Kurt Mehlhorn, and Karsten M. Borgwardt. Weisfeiler-Lehman Graph Kernels. *J. Mach. Learn. Res.*, 12:2539–2561, November 2011.

[Unger and Van Waarden, 2009] Brigitte Unger and Frans Van Waarden. How to dodge drowning in data? Rule- and risk-based anti money laundering policies compared. *Review of Law & Economics*, 5(2):953–985, 2009.

[United Nations Office on Drugs and Crime, 2011] United Nations Office on Drugs and Crime. *Estimating illicit financial flows resulting from drug trafficking and other transnational organized crimes*. Vienna, Austria, 2011.

[United Nations, 2000] United Nations. *United Nations Convention against Transnational Organized Crime*. November 2000.

[Wu *et al.*, 2019] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A Comprehensive Survey on Graph Neural Networks. *CoRR*, abs/1901.00596, 2019.

[Zhang *et al.*, 2018] Jiani Zhang, Xingjian Shi, Junyuan Xie, Hao Ma, Irwin King, and Dit-Yan Yeung. GaAN: Gated Attention Networks for Learning on Large and Spatiotemporal Graphs. *CoRR*, abs/1803.07294, 2018.