

NextGen AML: Distributed Deep Learning based Language Technologies to Augment Anti Money Laundering Investigation

Jingguang Han[†], Utsab Barman^{†§‡}, Jer Hayes[†], Jinhua Du^{†§}, Edward Burgin[†], Dadong Wan[†]

[†]Accenture Labs Dublin, Ireland

[§]ADAPT Centre, Dublin City University, Ireland

[‡]University College Dublin, Ireland

{jingguang.han, jeremiah.hayes, edward.burgin, dadong.wan}@accenture.com
utsab.barman@ucd.ie, jinhua.du@adaptcentre.ie

Abstract

Most of the current anti money laundering (AML) systems, using handcrafted rules, are heavily reliant on existing structured databases, which are not capable of effectively and efficiently identifying hidden and complex ML activities, especially those with dynamic and time-varying characteristics, resulting in a high percentage of false positives. Therefore, analysts¹ are engaged for further investigation which significantly increases human capital cost and processing time. To alleviate these issues, this paper presents a novel framework for the next generation AML by applying and visualizing deep learning-driven natural language processing (NLP) technologies in a distributed and scalable manner to augment AML monitoring and investigation. The proposed distributed framework performs *news and tweet sentiment analysis, entity recognition, relation extraction, entity linking* and *link analysis* on different data sources (e.g. news articles and tweets) to provide additional evidence to human investigators for final decision-making. Each NLP module is evaluated on a task-specific data set, and the overall experiments are performed on synthetic and real-world datasets. Feedback from AML practitioners suggests that our system can reduce approximately 30% time and cost compared to their previous manual approaches of AML investigation.

¹One of our banking clients has about 10,000 analysts globally to investigate over 400,000 red-alerted transactions weekly. With the data privacy and non-disclosure agreement, we are not entitled to release any clients' or entities' names.

1 Introduction

Money laundering (ML) is the process of transferring criminal and illegal proceeds into ostensibly legitimate assets, and it has long been considered as the world's third largest "industry". ML is often associated with terrorism financing, drug and human trafficking, so that it is a severe threat to the stability and security of the economy and politics globally. The International Monetary Fund (IMF) estimates that the aggregate size of ML in the world could be somewhere between 2% and 5% of the global Gross Domestic Product (GDP), which is equivalent to \$590 billion to \$3.2 trillion USD approximately. The annual growth rate of the volume of illicit funds traveling through ML channels is estimated as 2.7% (Jorisch, 2009).

AML is one of the long-standing challenges in the financial sector. A number of systematic frameworks have been proposed for AML systems (Gao and Xu, 2007; Gao and Ye, 2007; Gao and Xu, 2010; Gombiro and Jantjies, 2015) following a multi-stage procedure with data description and transaction evaluation approaches. Moreover, several techniques have been applied in this area: rule-based, link-analysis and risk-classification/scoring-based methods.

In the traditional AML system, financial transactions are consistently monitored using complex rules and thresholds to identify suspicious ML patterns and generate red alerts, e.g. unexpected high amounts or high frequent transactions. These rules are rigid and restrictive to ensure the blockage of ML transactions, resulting in a high amount of false positives (transactions blocked by mistake) (Pellegrina et al., 2009; Helmy et al., 2014). Consequently, a significant amount of human resources (reviewers/analysts) are engaged to approve or block such red-alerted transactions. This manual validation process is mainly a consultation

procedure which involves gathering intelligence information from different sources about the parties involved in a red-alerted transaction. Moreover, the manual process of gathering intelligence involves: (1) news search; (2) name screening; and (3) consultation of existing fraud database or criminal records to determine if any of the parties has a criminal record, fraud offense, or some other suspicious activities. The main drawbacks of most of current AML systems are: (1) a high volume of red ML alerts forces organizations to increase the amount of investigation costs in terms of time and human effort; (2) yet fewer ML alerts result in a limited number of checks, which heavily affects the recall of the system by allowing the suspicious transaction to pass through the compliance procedure (Gao et al., 2006); Several global banking organizations were heavily fined by AML regulators for ineffective AML practices (Martin, 2017).

We propose a novel framework that can drastically decrease the time and effort of *false positive validation* (FPV) for AML solutions. We harness multiple data sources, and apply deep learning-based NLP techniques within a distributed architecture to create a recommendation system to support human analysts for a more efficient and accurate decision-making. The salient contributions of this paper include: (1) Harnessing heterogeneous open data (e.g. social media, news articles and fraud bases) in AML compliance to make it more robust and updated; (2) Using different levels of sentiment analysis (SA) to identify negative evidence of a target entity; (3) Using name entity recognition (NER) and relation extraction (RE) to build the relation network of a target entity, and analyze the hidden and complex connections between the target entity and existing suspicious entities from the fraud bases; and (4) developing a distributed communication architecture for scaling the deployment of various NLP modules.

To the best of our knowledge, harnessing unstructured social and news content to facilitate the investigation and compliance has never been attempted in the recent literature of AML, nor any industrial systems.

2 System Architecture

The proposed novel AML framework follows a distributed architecture to integrate different deep learning-based NLP modules to facilitate the investigations of suspicious ML transactions. The

architecture of the system is presented in Figure 1.

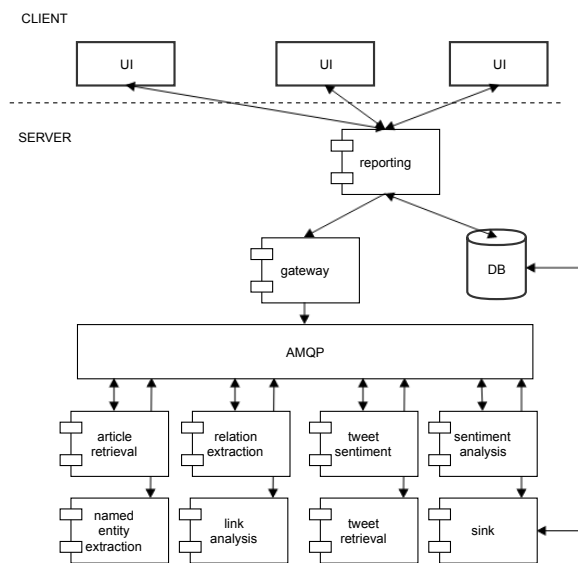


Figure 1: Architecture of the novel AML framework.

The system follows an extensible micro-service oriented distributed architecture where each component acts as a separate micro-service and communicates with others via an Advanced Message Queuing Protocol (AMQP) platform. AMQP is an open standard application layer protocol used for queuing and routing messages between the services in a secure and reliable way². Communications between the components are chained to achieve incremental data processing. This architecture also supports dynamic routing topologies.

In this framework, each AML component is independently configurable, deployable, scalable and replaceable which makes it flexible on where and how to run it. Thus, it can conveniently distribute the AML components over a cluster of machines where allocation of resources could be adjusted on demand depending on the needs regarding the processing time or memory consumption.

In Figure 1, the User Interfaces (UI) connect to the reporting service that provides the results information from the database (DB). The UI can also trigger the processing pipeline, calling the reporting service to make the gateway start processing data regarding new transactions.

The reporting service provides representational state transfer (REST) endpoints for the interaction between the UI and the system. This service allows analysts/experts to retrieve information re-

²<https://www.rabbitmq.com/specification.html>

lated to a specific transaction and the set of processed results that will lead to the human analyst to approve or block this transaction.

The information processing pipeline is a combination of different modules (micro-services) via different routings. Different routings encapsulate different customized functionalities for the end user. Using the route configuration file, the route is embedded in the message that is passed around, so the component knows where it should direct the message next.

In this architecture, the data layer is constructed using Cassandra, Neo4j and MySQL along with news and Twitter engines to govern, collect, manage and store various type of data (e.g. banking data and open data). The data layer maintains a bidirectional access with other modules in the system. Two types of data are handled in the data layer: banking and open data. Banking data refers to a wide variety of financial data, for example: data related to Know Your Customer (KYC), client profiles, customer accounts, and real-time transactions. On the other hand open data refers to financial news articles, social media, financial report, existing and open source fraud base etc.

3 Work-flow and Functionalities

The main UI of our system is shown in Figures 2–4. We will explicitly detail the functionality of each NLP module while following the work-flow of the system:

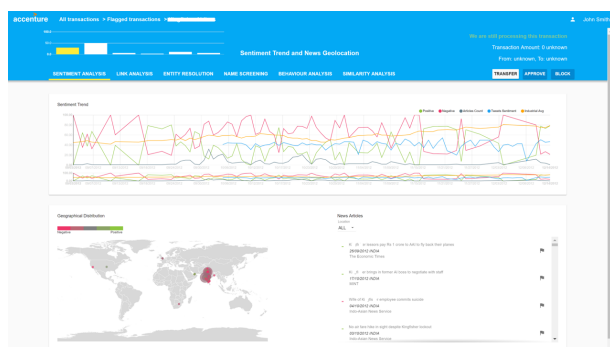


Figure 2: Sentiment Analysis of Target Entity

The system starts with transaction monitoring (TM) that identifies suspicious transaction using complex rules and generates red alerts. Suspicious transactions are flagged and reside in a queue for further investigation. Information such as name, location, account details of the parties (involved in a flagged) are extracted for further usage.

The name screening module filters out ML offenders or criminals have previous records of fraud, illegal activities (e.g. terrorism) and financial crimes by looking at a sanction list or ML cases from bank records and other fraud bases. The name screening employs fuzzy matching techniques as an initial filter like many off-the-self AML solutions. If the target entity is matched with some entities in the list, it will increase the probability of a ML transaction, so the system is triggered for more evidence gathering.




Figure 3: Knowledge Graph of Fraud Base: Each Node is an entity and edges are relations. Sensitive entities are blocked with white boxes according to data privacy agreement.

The target entity with its properties are stored in a knowledge graph (KG) which is constructed from a combination of publicly available fraud/ML bases³. Figure 3 shows part of the KG, where each node represents an entity and the edge represents a relation. Persons and organizations that occur in this KG and are red-alerted by the TM system, are strongly considered to be involved in illegal activities. Entity disambiguation is used when the the query of the target entity to the KG is to be made. If the target entity is matched, then we directly submit the evidence to human investigators. The system also considers human judgment in this process, mainly for two reasons: (1) Often fake identities, aliases are captured by analyzing the neighbors of the target entities (e.g. the associate entities remain same) (2) Often indirect connection to a suspicious entity validates the ML claim of a transaction.

Sentiment Analysis is a key module of our system that can be viewed as a sequential and as well

³<https://www.icij.org/investigations/panama-papers/>, <http://www.fatf-gafi.org/publications/> and <https://www.icij.org/investigations/paradise-papers/>

as a parallel component for evidence mining. At any given point of time the module starts fetching data from different sources, such as news sources and social media in order to find clues of the target entity, after that it analyses the sentiment of collected news and other data sources and projects the results into a geo-graph and a time-series chart to identify geo and temporal patterns. The goal of this module is to identify potential negative evidence involving crime, fraud and ML etc. As shown in Figure 2, the results are visualized with the time-stamp to reveal the trend. The idea behind viewing the sentiment over time is that *continuously growing* negative enduring sentiment for an entity indicates something suspicious and the necessity of in-depth investigation. In Figure 2, *Red* indicates the negative sentiment, *Green* represents the positive sentiment, and *Yellow* is the average sentiment of target entity’s competitors. *Blue* is the averaged sentiment from Twitter data. It can be seen that negative sentiment regarding to financial fraud is dominant for the target entity, and hence an in-depth investigation is necessary. The news are compiled and presented along with the graph so that the actual content can also be presented. Geographical distribution in terms of sentiment regarding the target entity is also visualized, which reflects the risks of the target entity involved in the ML activities in terms of location.

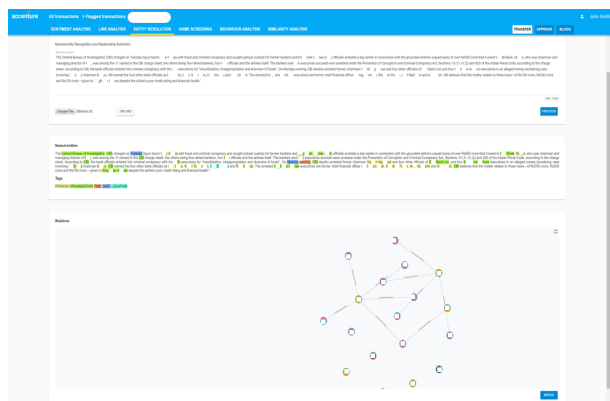


Figure 4: Knowledge Graph from News Article

The next part of our system is designed to perform entity resolution and linking through NER, RE, link analysis (LA). NER and RE are performed over collected unstructured negative news data to extract relational facts in real time and build an entity-specific KG regarding the target entity as shown in Figure 4. Subsequently, these entity-specific KGs are merged with the existing

knowledge graph by entity linking. In this process, an entity disambiguation technique is also used to find out the same entity with different mentions. Given the updated KG, we then carry out inspection and reasoning regarding the suspicious entity considering that a direct or indirect connection between a target entity and any number of existing ML/fraud entities. Direct and indirect connections can be a strong evidence implying that the investigated transaction might be illegal. This process not only serves evidence mining for a particular suspicious transaction but also enriches the existing KB continuously by adding a large number of entities that are collected from different sources.

Finally, for a suspicious transaction, confidence scores are generated along with evidences from each module: TM, fraud-base name screening, fraud KB matching, SA trend and entity disambiguation module. The scores are normalized within a range of 0.00 to 1.00 for each module and are presented to a user. Based on these scores, the user approves, blocks or transfers the transaction to higher authorities for action. The goal of our solution is not to replace human compliance personnels, but to augment their capability and to make a more accurate decision efficiently.

4 NLP Models

In this section, we introduce the modeling details for each NLP modules in our system.

4.1 Sentiment Analysis

Sentiment analysis is a task of identifying the polarity of sentiment (e.g. positive, neutral or negative) in a content at different levels, e.g. term, aspect, sentence and document.

For our AML task, we develop two different SA models, namely the document-level and sentence-level models. The former is a multi-channel convolutional neural network (CNN) based sentiment classifier (Kim, 2014) and used to process financial news articles; the latter is also a CNN based model (Tang et al., 2014; Deriu et al., 2016) for social media data.

One challenge for the document-level SA model is that annotated resources in the financial domain are hard to obtain. To solve this problem, we propose a voting scheme to label the training data as follows: Financial news are gathered using a list of keywords⁴ and news search APIs, afterwards they

⁴https://www3.nd.edu/~mcdonald/Word_

(headlines and the first paragraph of a news) go through some public sentiment APIs to generate the sentiment score. Finally, a voting mechanism is used to obtain the final result in terms of positive or negative sentiment for each document. Our document-level SA classifier, trained on automatic labeled 12,467 of such financial news articles, and it achieves 76.96% in terms of accuracy compared to a public sentiment API⁵ on the RT-polarity⁶ data set. With respect to the Twitter SA classifier, it is trained and evaluated on SemEval-2016 task 4⁷ data set, and achieves 63.10% in terms of accuracy, comparable to the best system (63.30%) in the shared task. Different from previous shared tasks, the SemEval-2016 task 4 is designed to estimate the percentage of tweets that are positive and the percentage of tweets that are negative in a given set of tweets about a topic. Thus, in this circumstance, this dataset is very helpful for us to verify our SA models for AML scenario because in one period there might have many tweets discussing a suspicious entity (e.g. an organisation).

4.2 Relation Extraction

Relation extraction involves the prediction of semantic relations between pairs of nominals or entities in a sentence (Bunescu and Mooney, 2005). We use the pipeline modeling methods for relation extraction tasks, i.e. recognising and disambiguating named entities (Wang et al., 2012) first, and then performing relation extraction on these recognised entities. NER is performed using a combined strategy: (1) Stanford NER Recognizer; (2) a neural NRE which is implemented using a LSTM-CRF framework (Lample et al., 2016); and (3) we combine the recognised named entities from these two models, and select out those that we want based on specific types. Seven different types of named entities are defined in our system, namely *Person*, *Organisation*, *Location*, *Date*, *Time*, *Money* and *Miscellaneous*. Given two entities and a sentence containing these two entities, our LSTM-based model predicts the relation between them. We evaluate it on SemEval 2010 task 8 (Hendrickx et al., 2010) data and it achieves

Lists.html

⁵We use <https://www.ibm.com/watson/alchemy-api.html> and it achieves 75.56% in terms of accuracy.

⁶<https://www.cs.cornell.edu/people/pabo/movie-review-data/rt-polaritydata.README.1.0.txt>

⁷Prediction of five-point scale polarity of a tweet.

80.62% in terms of macro-F1 measure.

Moreover, to handle multi-instance problem in the distant supervision RE (i.e. for one entity pair, there exists multiple instances containing this entity pair, where some instances are valid and some are noise), we develop an attentive RNN framework (Lin et al., 2016) with a word-level and a sentence-level attentions for relation prediction, where the word-level attention can learn lexical contexts and the sentence-level attention can select valid instances for relation prediction. We evaluate our model on the publicly available New York Time data set and achieves 88.00% accuracy in terms of P@100 measure.

5 Evaluation and Feedbacks from AML Practitioners

As discussed in above section, different validation and evaluation methods are applied to different NLP models, where the tweet SA model, news SA model or attentive RE model achieves or is comparable to the state-of-the-art in terms of accuracy. At present, the entire system is at piloting stage with our banking clients across US, Europe and Asia. It is currently being tested and evaluated by professional AML practitioners for AML and KYC investigations. From the feedbacks that we collected so far, the end users are optimistic on achieving the objective of reducing on average 30% of their time of investigating the red-alerted suspicious transactions and making a decision more efficiently. We have been invited to give keynote talks about different aspects (not the entire one) of this system at highly respected events, such as World Mobile Conference (WMC) 2018, Europe Financial Information Management Conference (FIMA) 2017 etc. Worth mentioning that some of our NLP models were also utilized by our clients in different domains. For instance, one of our diamond clients adopted our news and tweets sentiment analysis models for monitoring their brand reputation. Given the sensitivity of their business, they cannot release the performance metrics to us. However, their overall feedback and experience have been very positive.

6 Conclusions and Future Work

In this paper, we present a novel distributed framework of applying and visualizing different deep learning based NLP technologies to augment the anti money laundering investigation. Our system

is modularized and distributed which enables it to be deployed on scale and on demand. Each component is a micro-service which allows multiple instances of the same module to be created and deployed and used in tandem. By harnessing knowledge graph, sentiment analysis, name screening, named entity recognition, relation extraction, entity linking and link analysis, our system can provide different evidence extracted and analyzed from different data sources to facilitate human investigators. From the human evaluation and clients' feedback, our system can reduce by 30% in terms of human investigation effort.

In the future, we will (1) improve our models with more domain specific data, and fine tune the parameters; (2) scale and deploy the system on cloud-based servers for real-time processing of large volume of data; (3) tailor the solution and evaluate it in other domains such as KYC, and (4) adapt our system to multilingual use cases.

Acknowledgements

Many thanks to the reviewers for their insightful comments and suggestions. This work is supported by Accenture Labs Dublin, Enterprise Ireland under the Research Programme (Grant EI.IP20170626) and Science Foundation Ireland (SFI) Industry Fellowship Programme 2016 (Grant 16/IFB/4490).

References

- Razvan Bunescu and Raymond J Mooney. 2005. Subsequence kernels for relation extraction. In *In Proceedings of NIPS*, pages 171–178.
- Jan Deriu, Maurice Gonzenbach, Fatih Uzdilli, Aurelien Lucchi, Valeria De Luca, and Martin Jaggi. 2016. Swisscheese at semeval-2016 task 4: Sentiment classification using an ensemble of convolutional neural networks with distant supervision. In *Proceedings of the 10th International Workshop on Semantic Evaluation*, EPFL-CONF-229234, pages 1124–1128.
- Shijia Gao and Dongming Xu. 2007. Conceptual modeling and development of an intelligent agent-assisted decision support system for anti-money laundering. *Expert Systems with Applications*.
- Shijia Gao and Dongming Xu. 2010. Real-time exception management decision model (rtmdm): applications in intelligent agent-assisted decision support in logistics and anti-money laundering domains. *International Conference on System Sciences*.
- Shijia Gao, Dongming Xu, Huaqing Wang, and Yingfeng Wang. 2006. Intelligent anti-money laundering system. *International Conference on Service Operations and Logistics, and Informatics*.
- Zengan Gao and Mao Ye. 2007. A framework for data mining-based anti-money laundering research. *Journal of Money Laundering Control*.
- Cross Gombiro and Mmaki Jantjies. 2015. A conceptual framework for detecting financial crime in mobile money transactions. *Journal of Governance and Regulation*.
- T.H. Helmy, M. Zaki, T. Salah, and B.K. Tarek. 2014. Design of a monitor for detecting money laundering and terrorist financing. *Journal of Theoretical and Applied Information Technology*, 1(1):1–11.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2010. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Proceedings of the Workshop on Semantic Evaluations: Recent Achievements and Future Directions*, pages 94–99.
- Avi Jorisch. 2009. *Tainted Money: Are we losing the war on money laundering and terrorism financing?* Red Cell Intelligence Group.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). *CoRR*, abs/1603.01360.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th ACL*, pages 2124–2133.
- Ben Martin. 2017. Deutsche bank hit with pounds 500m money laundering fines. <http://www.telegraph.co.uk/business/2017/01/31/deutsche-bank-hit-500m-money-laundering-fines/>.
- D. Pellegrina, L. Donato, and M. Donato. 2009. The risk based approach in the new european anti-money laundering legislation: a law and economics view. *Computing Reviews*, 5(2):290–317.
- Duyu Tang, Furu Wei, Nan Yang, Ming Zhou, Ting Liu, and Bing Qin. 2014. Learning sentiment-specific word embedding for twitter sentiment classification. In *Proceedings of the 52nd ACL*, pages 1555–1565.
- Longyue Wang, Shuo Li, Derek F Wong, and Lidia S Chao. 2012. A joint chinese named entity recognition and disambiguation system. In *Proceedings of the Second CIPS-SIGHAN Joint Conference on Chinese Language Processing*, pages 146–151.