Digital Finance
https://doi.org/10.1007/s42521-020-00023-1

ORIGINAL ARTICLE

# Artificial intelligence for anti-money laundering: a review and extension

**Jingguang Han[1] · Yuyun Huang[2] · Sha Liu[2] · Kieran Towey[3]**

Received: 8 October 2019 / Accepted: 16 June 2020


## Abstract

This paper surveys the existing academic literature on artificial intelligence (AI) technologies for anti-money laundering (AML). We review the state-of-the-art AI methods for AML and extend the discussion by proposing a framework that utilizes advanced natural language processing and deep-learning techniques to facilitate next-generation AML technologies. Our framework utilizes unstructured external information to assist domain experts, aiming to decrease the workload for the human investigator. We bridge the gap between the current AML methods and state-of-the art AI, highlighting new trends and directions in AI that can be used to develop the AML pipeline into a robust, scalable solution with a reduced false positive rate and high adaptability.

**Keywords** Anti-money-laundering · Artificial intelligence · Natural language processing · Deep learning

**JEl Classification** G21 · G23 · C44 · C45

✉ Sha Liu
sha.liu@ucd.ie

Jingguang Han
hanf@tcd.ie

Yuyun Huang
yuyun.huang@ucd.ie

Kieran Towey
kieran.towey@kpmg.ie

[1] Vanke Service Research, Vanke Buidling, Meilin Road, Futian District, Shenzhen, China

[2] Michael Smurfit Graduate Business School, University College Dublin, Dublin, Ireland

[3] KPMG, 6 Ballynakelly View, Newcastle, Co. Dublin, Ireland

Published online: 25 June 2020

## 1 Introduction

Money laundering is legally defined as "transferring illegally obtained money through legitimate people or accounts so that its original source cannot be traced" (Black's Law Dictionary 2009: 1097). The International Monetary Fund (IMF) estimates the aggregated size of worldwide money laundering as approximately $3.2 trillion, or 3% of the global GDP (Jorisch 2009). The profits of money laundering are often used to finance crimes, including terrorism, human trafficking, drug trafficking, and illegal arms sales (Jorisch 2009).

Anti-money laundering (AML) systems are implemented by financial institutions such as banks and other institutions that provide credit, in an effort to combat money laundering by identifying money laundering risks, potential money launderers, and money laundering transactions (Unger and Waarden 2009). The manner in which financial institutions run their business, the risks they take, and the policies they implement (or not implement) must pass the scrutiny of third parties, including customers, shareholders, governments, and regulators (Parkman 2012). Falling short of the required AML standards is a form of corporate wrongdoing, with these financial institutions facing reputational risk; this is defined by the US Federal Reserve Board (2017) as "the potential that negative publicity regarding an institution's business practices, whether true or not, will cause a decline in the customer base, costly litigation, or revenue reductions" (pp. 6.5–6.6). The consequence could be a significant *reputational cost*, which can be estimated—for publicly listed firms—as the difference between the loss of the perpetrating firms' (e.g., the banks') share market valuations and their direct regulatory fines, legal fines, and other costs. The loss of market value can amount to many times that of the direct legal fines and costs (see Barnett 2014; Bernile and Gregg 2009; Johnson et al. 2014; Karpoff et al. 2008; Murphy et al. 2009; Yu et al. 2008). Despite current efforts, several multinational financial institutions have been heavily fined by AML regulators for ineffective AML practices in recent years (Viswanatha and Wolf 2012; Titcomb 2014; Thompson and Perez 2017; Martin 2017; Irish Examiner 2016). The General Data Protection Regulations (GDPR), which have been in effect since May 2018, further stress the importance of financial institutions' AML efforts.

This paper surveys the existing academic literature on artificial intelligence (AI) for AML. Mei et al. (2014) found that in the preceding 20 years (1994–2014), US-based researchers alone published 97 papers on AML, the largest output worldwide. Chinese researchers published 42 papers. However, until now, few papers have surveyed existing studies on the application of AI or data mining to AML. This paper is among the first to systematically review AI methods for AML.

We begin by reviewing the current policies and frameworks for identifying money launderers and money laundering transactions. We then survey the existing studies on automatic AML, to show the full scope of AML techniques. Currently, deficits exist between financial institutions' AML systems and the state-of-the-art AI solutions; their current AML systems operate via a combination of

human expertise and machine automation. These methods often incorporate AI or data-mining techniques; however, there remains a strong dependence on human auditors. Financial institutions' AI systems tend to be simplistic and rule-based; a transaction will be flagged as suspicious and require a human-conducted review to determine if it fails to pass a set of rules outlined by the governing authorities. Rule-based systems lead to an unmanageable number of transactions being flagged as suspicious, requiring a large amount of time and money to be spent on reviewing legitimate transactions (Gallo and Juckes 2005). Institutions must strike a balance between the rigorous overview of transactions and the approval of legitimate transactions in a timely manner. The introduction of AI enhances and facilitates the overall decision-making process whilst remaining compliant with policies such as the new GDPR. AI can minimize the number of transactions falsely labelled as suspicious, achieve a demonstrable quality of compliance with regulatory expectations, and improve the productivity of the operational resources. In this paper, we detail the technical challenges that may be faced by the AML community and how these can be mitigated by implementing AI methods.

Our most significant contribution is that we offer a unique perspective of AML by addressing how AI can push AML forward. We focus on how AI can immediately improve AML systems and help guide AML solution designs in the future. The AI methods currently being published in the AML literature are outdated and have yet to adopt some of the most impactful AI techniques, such as deep learning. In this paper, we go beyond the existing knowledge and propose next-generation AI methods and systems (incorporating deep learning and natural language processing) to address the challenges of AML operations. In summary, we bridge the gap between the current AML methods and state-of-the art AI, highlighting new trends and directions in AI that will develop the AML pipeline into a robust and scalable solution with fewer false positives and a high adaptability. Furthermore, we aim to publish a data corpus for AML research. It would be the first corpus in the field that parses texts into a more comprehensive and structured form for easier interpretation of advanced analytics and AI algorithms.

The remainder of this paper is structured as follows. In Sect. 2, we review the current AML policies (including rule-based and risk-based policies) and the new GDPR. In Sect. 3, we review a range of AML studies and chart the current academic landscape of automatic and smart AML. We look in detail at papers that propose AML frameworks, as well as those that discuss rule-based methods, link analysis, outlier detection, and classification methods for the automatic detection of money laundering activities. We discuss the limitations of each method and provide possible solutions to common concerns and constraints. We conclude this section by highlighting general AML survey papers for the reader's reference. In Sect. 4, we propose a next-generation AML system using a multichannel fusion of state-of-the-art deep learning and natural language processing techniques. We detail the features of our system, describing how they help identify frauds and communicate them to investigators. Section 5 concludes the paper.

## 2 Anti-money laundering policies

The Financial Action Task Force (FATF) was created in 1989 to combat money laundering and terrorist financing. It is an intergovernmental agency consisting of 35 member jurisdictions and two regional organizations. The FATF requires that all AML and Counter Terrorist Financing (CTF) programs include specific data analysis and reports. They further require that an institution be able to identify and verify their clients, commonly referred to as the "know your customer" requirement. This prohibits anonymous accounts and fictitious account holder names, and demands that institutions implement preventative measures when dealing with correspondent and shell banks. Another requirement is that banks must keep records of all transactions for a minimum of 5 years. The data need to include the names of customers and/or beneficiaries, their addresses, the nature of the transactions, the dates of transactions, the types of currency, the amounts of currency, the types of account, and the identifying numbers of any account used. The FATF requires two types of reporting: suspicious transaction reports (STRs), which are filed with the national financial intelligence unit, and currency transaction reports (CTRs), which report transactions above a certain amount.

### 2.1 Rule-based policy

The "Forty Recommendations on Money Laundering" were released in 1990 as a basic framework for preventing, detecting, and suppressing illicit financing. After the September 11, 2001 terrorist attacks in New York City, the "Nine Special Recommendations on Terrorist Financing" were released to focus efforts on CTF. Such rule-based policies are clear and transparent, making them easy for financial institutions to implement and show compliance. Current AML systems use these rules in their automatic layer, to decide if a transaction is suspicious. However, the clarity of these rules makes it harder for fraud to be detected because transactions can be designed to bypass the rules. The rules also lead to an over-reporting of suspicious behavior, which is expensive and time consuming to filter through.

### 2.2 Risk-based policy

A risk-based policy gives private actors more discretion on what to report, by having vaguer reporting criteria. This puts the decision in the hands of private businesses, making them responsible for the success of reporting. However, this has made over-reporting worse in some countries. It is a higher-risk method that can be arbitrary or haphazard, causing banks to lose customers or be fined for under-reporting.

Nevertheless, many financial organizations have adopted risk-based systems (Helmy et al. 2014). They overcome the limitations of rule-based solutions by assessing client and transaction risks, and by identifying outlier behavior.

### 2.3 Enforcement and detection

The FATF calls for the criminalization of money laundering. Not all countries have done this, and the law enforcement is not and cannot be uniform across nations; thus, different nations have adopted a subset of the following enforcement methods, with varying degrees of success: (1) monitoring imports and exports between countries; (2) requiring charities to file annual documentation, including financial statements; (3) blacklisting countries and suspected terrorist financiers; (4) sanctions; (5) prison sentences, penalties, and fines; (6) demanding bank transparency; and (7) heavily fining banks that do not sufficiently investigate money laundering fraud (Viswanatha and Wolf 2012; Titcomb 2014; Thompson and Perez 2017; Martin 2017; Irish Examiner 2016).

### 2.4 General data protection regulation (GDPR)

Recently, the European Union (EU) has adopted some important regulations (formally known as GDPR) regarding the collection, storage, and use of personal information; these took effect as law across the EU from May 2018 and replaced the EU's 1995 Data Protection Directive (DPD). In terms of scope, the new regulation is uniformly applicable to industries, EU organizations, and organizations that deal with EU data. The data-driven regulations focus attention on some specific issues, including the ownership of data, transparency, explainability, and the trust capacity of algorithms that are trained or built with such data. A detailed analysis of these regulations can be found in (Goodman and Flaxman 2016; Kamarinou et al. 2016; Kuner et al. 2017). To summarize, GDPR stipulates that data-driven automated systems—including AML systems—must adapt the following during implementation: (1) legal data processing and data ownership, (2) explanatory frameworks for the data and algorithms, and (3) ethical compliance.

## 3 State-of-the-art AML with AI

A number of frameworks have been proposed for AML systems (Gao and Xu 2007,2010; Kolhatkar et al. 2014; Gombiro and Jantjies 2015; Lai 2018; Weber et al. 2018). Typically, these follow the multistage approach described below; they begin with data descriptions and then progress to smart transaction-evaluation approaches. Gao et al. (2006) proposed a system to determine if a transaction is a high-risk transaction, using Simon's model of decision-making (Gao and Xu 2007). Gao et al. (2006) presented an intelligent AML system that employs human agents, arguing that these are necessary for the system to be able to learn and adapt. Gao and Ye (2007) recommended using link analysis, unsupervised techniques, and entity extraction to predict money laundering. A major shortcoming of these papers is that—although they recommend the use of AI—they do not give any details as to

how the technology should be implemented and designed, or what it would be able to contribute. We intend to fill the gap left by these papers—that is, their lack of technical details—in Sect. 4.

### 3.1 Money laundering phases

*Placement*, *layering,* and *integration* are the three phases in money laundering schemes. Proceeds from criminal activities enter the placement phase, where they are converted into monetary instruments or otherwise deposited in a financial institution (or both). Layering refers to the transfer of funds to other financial institutions or individuals via wire transfers, checks, money orders, or other methods. In the final phase of integration, funds are used to purchase legitimate assets or to continue financing criminalized enterprises. Here, illegally obtained money becomes part of the legitimate economy.

AI approaches may be applied to identify money laundering activities in each of the above three phases. Common machine learning methods such as support vector machines (SVMs) and random forest (RF) can be used to classify fraud transactions, using large, annotated bank datasets (Tang and Yin 2005). These data-driven approaches are normally used for the placement and layering phases because the transaction data is monitored by the bank. The final phase of integration is difficult to detect because funds have passed fraud-detection mechanisms. At this stage, advanced AI methods—for example, entity relationship extraction from large social media and news data—could be applied to AML.

Once suspicious transactions have been identified by rules or machine learning-based systems, a fraud investigator will become involved in the following analytical procedures. The workload of a human investigator largely depends on the number of fraud transactions reported. Natural language processing (NLP) approaches of entity and relationship analysis can help relieve the work burden by providing the human experts with a score and link relationship visualization based on news data (e.g., the banks' news database and traditional or social media news sources) concerning the potential fraud entity, using NLP and knowledge-based technologies.

In the following subsections, we review some commonly used AML solutions.

### 3.2 AML industry solutions

At present, the typical AML workflow in industry is a linear pipeline that connects a data source to a rule-based system. Analysts then incorporate their own research to determine if transactions are legitimate or fraudulent. A specific multistage process is followed. First, AML frameworks collect and process data. Second, they screen and monitor transactions. If a transaction is found to be suspicious, it is flagged, and a human analyst will decide if the flagged transaction is fraudulent. Generally speaking, AML frameworks can be decomposed into four layers. The first layer is the *Data Layer*, in which the collection, management, and storage of relevant data occurs. This includes both the internal data from the financial institution and external data from sources such as regulatory agencies, authorities, and watch-lists. The

second layer, the *Screening and Monitoring Layer*, screens transactions and clients for suspicious activity. This layer has been mostly automated by financial institutions into a multistage procedure often based on rules or risk analysis. If a suspicious activity is found, it is passed on to the *Alert and Event Layer* for further inspection. This process includes augmentation of the data with historical transaction information and necessary evidence, to review the flagged transaction. Harnessing social media and web content to acquire information for investigation is underdeveloped in current AML systems. Consequently, auditors are under-resourced, increasing the inaccuracy of auditor decisions and the time required to inspect each transaction. The decision to block or approve a transaction is made by a human analyst in the *Operational Layer*.

### 3.2.1 Data layer

The first layer in the AML framework is the *Data Layer*, in which data are collected, managed, and stored by various submodules and agents. This layer maintains bidirectional access with other layers and handles both internal and external data.

Internal data refers to a wide variety of data that are identified and processed internally by the different components of a system. Some data sources are accessed and obtained directly; for instance, client profiles, customer accounts, and real-time transactions. These are used for client profile assessment, transaction risk measurement, and behavior diagnosis. The outputs of different analytics engines, insights from the analysis, and the histories of previous blocked transactions are also handled in this layer. These data are used to validate the final decisions made when evaluating transactions throughout the system. External data are data collected from sources outside of the financial institution; these can include regulatory agencies, government authorities, international standards, legislation, sanctions, and watchlists. Often, social media and news portals are considered as external data sources; however, they are currently underutilized in AML solutions. From a technological point of view, traditional systems suffer from architectural deficiencies such as data quality, data management, and data governance issues. Big-data technology and distributed data processing have not yet been widely implemented in the AML community.

Large quantities of heterogeneously formatted data are used in AML systems. The *Data Layer* is typically maintained by an enterprise data hub that incorporates the technologies relevant for efficient processing. For example, Hadoop is used for parallel processing and data collection (White 2009), Solr is used for searching, and Mahout or Spark are used to model and produce analytics (Owen et al. 2011; Zaharia et al. 2010). Different databases (relational and linear) are used to store the raw and processed data, as well as the analytical results. For better understanding, we divide this layer into the following components, based on their purpose:

*Collection agent*: the collection component deals with internal and external data collection. Generally, this component has a distributed data collection and processing framework. Technologies such as Hadoop, Kafka, and Storm are used here.

*Processing agent*: data processing is an important but cumbersome task. Because the data originate from multiple sources, they are collected and generated with

different standards. All incoming data need to be standardized for later use. The data are often enriched with meta-data generation and linking techniques, and they are formatted and compressed using information-retrieval tools (e.g., Solr) and NLP techniques, including tokenizers.

*Insight agent*: this component generates insights from the data using several data analysis techniques. For example, similar transactions and customers are clustered together according to their profiles. Links between customers and transactions are created using data analytics tools such as Mahout. Anomalous transactions can also be identified using the customer or organization profiles. This information is used in the subsequent layer to determine if a transaction should be regarded as suspicious.

*Storage agent*: data storage is a critical component of this layer. Because data are gathered, generated, and processed in this layer, they also need to be stored in relational (e.g., Oracle) or linear databases (e.g., Cassandra). Often, several data-management frameworks are incorporated in this layer.

*Security agent*: data security is one of the primary concerns of financial institutions. Sensitive financial data—such as credit card and account information—must be protected in the database and in the pipeline of the banking system. Moreover, organizations must fulfil the requirements of global data-protection policies. The *Security Agent* serves these requirements; the main objectives of this layer are the prevention of data breaches (by safeguarding the sensitive financial data) and adjusting security measures to match emerging forms of cyber-attack. Technologies such as secure key management, access controls, data-access monitoring, firewalls, and advanced encryption techniques are used by this agent.

### 3.2.2 Screening and monitoring layer

Screening and monitoring comprise the second layer in the AML system. Screening occurs before a transaction is executed, it consists of name and transaction screening. The monitoring process is performed continuously, it surveys the transactions and client profiles with the help of analytical models. The components in this layer operate in a collaborative framework that involves several tools. They maintain a bi-directional connectivity with the *Data Layer* for data retrieval and post-operational storage. For a more detailed understanding, we divide these components according to their specific tasks:

*Transaction-screening module*: this module operates before a transaction is executed; it is used to comply with sanctions. Generally, the set of rules defined by the Wolfsberg Statement of AML Screening Monitoring and Searching is followed worldwide in real-time transaction screening. The transaction-screening module maintains connections with the *Data Layer*, to receive external data for filtration. Applications that provide transaction-screening services include Actimize and MANTAS.

*Name-screening module*: this module is used to identify payments relating to people or organizations that have been identified as potential money launderers by regulatory authorities. The checks are performed continuously and in real-time; this requires the module to maintain connections to the *Data Layer*. Advanced matching algorithms (including entity resolution, discussed in Sect. 4.1.1) are critical in

this step. Several organizations deliver quality name-screening services to financial corporations, including Compliance Link of Accuity, Oracle Watch-list Screening of Oracle, and LexisNexis Bridger Insight XG of LexisNexis Risk Solutions.

*Transaction monitoring module*: this module identifies suspicious transaction patterns and completes a suspicious activity report or an STR. Different data mining, AI, and visualization techniques are implemented by this module, including link analysis (discussed in Sect. 3.4) and outlier detection (Sect. 3.5). Alongside its own analysis engine, this module incorporates the results of the screening modules and is connected with the *Data Layer*, to retrieve information and store reports.

*Client profile-monitoring module*: this module analyzes a client's account to provide an overview of the client profile. It also maintains a bi-directional connection with the *Data Layer* to receive client data and store analytical results. The component operates in collaboration with other modules in the *Screening and Monitoring Layer*; however, it specifically focuses on certain activities, such as alerting high-risk countries, analyzing financial connections and business relationships, and understanding political affiliations. Often, client profiles are compared with a group of potential or acknowledged money launderers, to obtain a similarity metric. Link, pattern, and risk analysis are common techniques used in this module.

Rule-based systems depend upon human-defined rules and thresholds, which are easy for launderers to understand and thus avoid violating. Furthermore, if the rule-based system is very strict, a large number of transactions will be falsely labeled as suspicious (Lucia and Donato 2009; Helmy et al. 2014), leading to a substantial number of manual inspections. On the other hand, being insufficiently strict and only flagging transactions if they exceed a high threshold results in a small number of checks and the acceptance of too many illicit transactions (Gao et al. 2006). Thus, most of the rule-driven AML solutions are incapable of adequately handling large amounts of transactional and financial data, making them impractical on the scales experienced by banks. They are also unable to generalize or automatically adapt to new crime patterns, because the rules are defined in advance. The results of the screening and monitoring modules form the basis of alert generation. Transactions considered as suspicious are flagged for further processing by the alert and operational layers.

### 3.2.3 Alert and event layer

The *Alert and Event Layer* raises an alert if a suspected transaction needs to be reviewed by a human evaluator. The large numbers of transactions to be reviewed—and the sparse supporting data provided for the investigation—increase the time required to inspect each transaction. Financial corporations are applying state-of-the-art statistical and data-related technologies in their AML approaches, to reduce both the risks and costs of manual inspection. This layer bases its decision on the historical data of previous decisions and through comparisons to similar transactions and decisions. If a transaction is flagged, the layer augments it with additional data for the evaluator. This includes a history of decisions on similar transactions, the risk scores computed by previous layers, and the priority of clearing the transaction.

Pending transactions—for which no decisions have been made—are stored in a backlog. The high number of alerts generated by the previous layers often leads to a large backlog. Alerts, produced and prioritized by this layer, finally appear in the *Operational Layer* for manual intervention to permit, block, or reject the transaction. This layer suffers from a large number of false positives—legitimate transactions that are falsely flagged as suspicious. This leads to a large backlog and often an overwhelming number of tasks for the human evaluator.

#### 3.2.4 Operational layer

In the final layer, human agents finalize the decision to block, release, or queue a transaction, based on the data received from the previous layers. It is a legal requirement that the final decision concerning a transaction be made by a human agent. The previous layers serve to monitor all transactions and flag potentially fraudulent transactions. However, the final decision is made by a person using the information provided by the preceding steps in the process.

The human agents use a range of techniques to compile and visualize supplemental information concerning the suspicious transaction. This includes querying the World Wide Web (WWW) and LexisNexis for information about the entities affiliated with the transaction and visualizing the connections the entity has using link analysis. In Sect. 4.1.1, we propose additional AI methods to improve this layer, including sentiment analysis, entity resolution, and knowledge graphs.

Rule-based industry solutions dominate the set of current methods applied by financial institutions. In these solutions, predefined rules (e.g., transaction thresholds) are applied to an incoming transaction, to determine whether or not it is suspicious. Though simplistic and easy to fool, using rules is commonplace because they make it easy to demonstrate compliance with regulations. More intricate implementations of rule-based systems exist; for instance, Rajput et al. (2014) categorized specific accounts and transactions by constructing a reasoning based on an ontology, then queried that ontology with new transactions. Moustafa et al. (2015) used hard-coded rules, searching for transactions from suspected countries, people, organizations, and accounts. They also searched for transactions that exceeded a certain threshold. Furthermore, they employed link analysis as a visualization tool, to identify indirect connections to suspicious entities. They also looked for cycles in the linked graph, to spot money laundering occurring across multiple transactions.

### 3.3 Network analysis for AML

Network analysis is another method for identifying money laundering activities. Network analysis in the AML research field typically refers to those studies using relational data to locate direct and hidden connections with a money laundering node. One of the initial methods of network analysis was centrality evaluation, which was used to determine which node was the most important in a network (Bavelas 1950). Common network analysis systems contain the following variables: degree of centrality, authoritativeness, betweenness centrality, closeness centrality,

hubness, and page rank (Hanneman and Riddle 2005). Degree centrality indicates the number of direct connections of one node/vertex in the network. Authoritativeness is the degree to which one node points to another node via important hubs. Page rank measures the fraction of time spent on one vertex over all other vertices; it reflects the importance of a node.

Betweenness centrality is defined thus: for a graph $G = (V, E)$, the betweenness centrality of a vertex $v$ is expressed as $C_{\mathrm{B}}(v) = \sum \frac{\delta_{\mathrm{st}}(v)}{\delta_{\mathrm{st}}}$, where $\delta_{\mathrm{st}}$ represents all short paths from $s$ to $t$ and $\delta_{\mathrm{st}}(v)$ represents the $s$ to $t$ short paths that contain vertex $v$.

Closeness centrality measures how close a vertex is to others. The closeness centrality of vertex $v$ in a graph $G = (V, E)$ is formulated as $C_{\mathrm{C}}(v) = \frac{1}{\sum_{u \neq v \in V} L_{\mathrm{uv}}}$, where $L_{\mathrm{uv}}$ is the length of the shortest path from $u$ to $v$.

Drezewski et al. (2015) applied network analysis components to construct and analyze social networks using bank statements and National Court Register data for money laundering cases. During the network analysis procedure, they assigned roles to nodes in the network, measured the roles' connections, attempted to determine the entities' mutual proximity, and compared this information to the external role information (e.g., bank statements and the National Court Register) assigned to nodes.

Colladon and Remondi (2017a, b) built several networks to work collectively, including a transactions, economic sector, geographical area, and tacit link network to prevent money laundering. They used the actual, 19-month data of a factoring company that mainly operates in Italy. They found that network metrics were extremely useful in fraud-risk assessment.

### 3.4 Link analysis

One approach to identifying money laundering is to define a linked graph over entities. Relationships between subjects (represented as nodes) can be identified by the transactions that connect them (represented as links). This approach has a rich history in the literature, including the studies of Goldberg and Wong (1998) and Senator et al. (1995), in which linkage and case-based reasoning were used to visualize and analyze money laundering. These ideas have also been applied to mobile money laundering in Lopez-Rojas and Axelsson (2012a). They created synthetically linked graphs (similar to social networks), which they used to visualize and detect certain connections. Zhang et al. (2003) considered a situation in which no explicit links were observable between entities in a linked graph. They created communities—based on yet-to-be determined relationships—and used correlations as the attributes, to form new links. They used 7668 free-text documents regarding a real money laundering case, scanned the documents using optical character recognition, tagged the key entities, and created an extensible markup language file containing details such as the person's name, the organization, transaction time, transaction location, amount, and so on. They tested their link-creation algorithm on this data.

Link analysis is a useful tool for representing the connections between entities (e.g., subjects, organizations, bank accounts). In Sect. 4.1.1, we detail state-of-the-art machine learning methods that can enhance the links in a graph (relation

extraction, entity resolution) and the information captured by the graph (knowledge graph).

### 3.5 Outlier detection

A natural way to frame the AI and data-mining tasks of laundering or fraud detection is through outlier detection. In this method, one defines what a normal or inlier transaction would appear as for a subject and then detect any transaction that is sufficiently different to be considered as an outlier. A peer group is defined, to capture the typical spending habits of a customer. Clustering is a standard method of defining peer groups; next, a distance is computed between the incoming transactions and peer groups to detect outlier behavior (Hand and Weston 2008; Zengan 2009; Raza and Haider 2010; Kannan and Somasundaram 2017).

For example, Le-Khac et al. (2009) used k-means clustering to cluster data, then used a new transaction's distance from the clusters to identify outliers, using information provided by a bank. Larik and Haider (2010) also employed clustering to define the normal behavior. They ranked the incoming transactions according to their deviation from the clusters, using approximately 8.2 million real transactions. Liu et al. (2008) found suspicious transaction *sequences*, using individual account history and information from other similar accounts. They used data provided by a Chinese financial institution and calculated the similarity between new transactions and high-risk transactions, to rank suspicion.

### 3.6 Risk classification/scoring

A number of academic papers present the *Screening and Monitoring Layer* as a classifier that determines if a transaction is suspicious. Due to the lack of real data, many groups have worked with simulated datasets. Furthermore, because financial institutions are rarely informed if a transaction was determined to be money laundering (that decision being made by government officials), the existing methods often use synthesized fraudulent data. For example, Tang and Yin (2005) trained an SVM to predict suspicious transactions from real bank data to which they had appended synthesized suspicious data. Lopez-Rojas and Axelsson (2012b) tested a number of classification techniques—including random forest, decision trees, naive Bayes, and decision table—for predicting money laundering in mobile applications. They used the customer ID, profile, date of transaction, type of transaction, amount transferred, location, and customer age to represent each transaction. They created 486,977 synthetic transactions, 6,006 of which were labeled suspicious. They found that decision trees achieved the highest performance. A decision tree is the representation of a rule-based system that categorizes transactions as fraudulent if they exceed certain thresholds. Machine learning and data-mining methods have been applied to discriminate fraudulent transactions and to predict whether new transactions are fraudulent. Kingdon (2004) predicted outliers by training an SVM to identify unusualness, based on features that represent the activities of users. This process is very similar to the outlier-detection methods discussed above; however, it differs in that it

trains a model to predict unusualness rather than matching transactions to previous instances using clustering.

Some authors have been able to experiment with real data provided by government sources. For example, Paula et al. (2016) were given Brazilian import and export data by the Secretariat of Federal Revenue of Brazil. They trained an unsupervised auto-encoder to identify fraudulent exports. They verified their results using people but did not report any results. Wang and Yang (2007) applied AML techniques on a Chinese market and argued that a market-specific model is required. They trained a decision tree on the data of 28 customers who were represented by attributes such as the industry they worked in, their location, business size, and deposits. Their decision tree determined whether a transaction was of low, medium, or high risk. Helmy et al. (2014) ran a case study on a single donation that masked a money laundering transaction. They used a rule and risk-based system to identify whether the transactions were certain, likely, possible, or rare. They trained a finite-state machine to represent money laundering scenarios and detected cycles within graphs to identify suspicious links between entities. Lv et al. (2008) were given eight months' worth of data from a financial institution holding 1 million records from 6000 accounts. They used a triple-layer neural network to predict if transactions were suspicious or not; they used one hidden layer, whereas current state-of-the art neural-network methods use hundreds of layers. Le-Khac and Kechadi (2010) used a case study from an international investment bank. They took information regarding six funds from 10,000 customers over 14 years, extracting features to describe the funds and clustering them to create groups representing types of activity. They then found outliers based on the distance of a transaction from the clusters. Khan and Haider (2013) also used real data and built a Bayesian network, to predict if a transaction was suspicious based on certain rules that they defined. Colladon and Remondi (2017a, b) predicted and assessed the risk of clients, using real data from a factoring company. They created a linked transaction network and applied social network metrics to assess risk profiles.

### 3.7 Graph learning for AML

A node can represent a single account which itself also forms another graph in massive transaction-graph data. Meanwhile, financial institutions process millions of transactions per second. The main challenges in graph learning for AML are the graph learning/parsing speed and graph size. A preliminary work by Weber et al. (2018) focused on a faster graph-learning technique for AML. Fast-graph learning utilizes Fast-Graph Convolutional Networks (fast GCN), it dramatically increases training speeds compared with conventional GCNs.

### 3.8 Shortcomings of current solutions

The definition of money laundering itself is a reoccurring problem for policy makers and for AML system designs. Because there is no single pattern that identifies fraud, money laundering can be easily confused with legitimate transactions. The patterns

of frauds are constantly changing, making it hard for rule-based systems and policies to keep pace. These difficulties lead to institutions having to choose between the efficiency and effectiveness of their AML system. An efficient system that quickly determines a fraud relies less on human analysts and presents a higher risk of overlooking fraudulent transactions. A lower-risk system is more secure because the majority of transactions will be screened in depth; however, this is costly to both the analysts and the financial institution. The trade-off between risk and cost must be taken into consideration by individual institutions when they design their systems. AI is a key method for addressing screening and adaptability issues.

Research into automatic AML has been limited by data access and technical issues. It is important for systems to be trained using real-world data, instead of simulated data. At present, there are no open-source data for money laundering research, due to the importance of maintaining client privacy. Therefore, data need to be provided by private institutions; this is a difficult task because releasing client data can compromise an institution's reputation and may not comply with data privacy governance. Institutions struggle to process their data internally owing to the large quantities of data they collect, much of which is noisy. We detail data-management solutions in Sect. 4.3.1.

One of the most significant issues for financial institutions to overcome in implementing AML is the distribution, storage, and processing of their data. In most institutions, the majority of transactions are legitimate and should be accepted. Any transactions that are flagged as fraudulent by analysts are sent to the authorities for further evaluation. The final decision upon whether a transaction is truly fraudulent is not necessarily reported to the financial institution. Therefore, financial institutions obtain a large number of suspicious transactions, a subset of suspicious transactions that their analysts flag as fraudulent, but no feedback on the accuracy of their decisions from the authorities. This means that the data that financial institutions can provide researchers is noisy and often comes without a label of "fraudulent" or "legitimate;" this makes it difficult to construct models for predicting fraudulent transactions. It also makes it difficult to evaluate the accuracy of existing methods through comparison with each other. Without a true positive and true negative class, it is impossible to state whether a method works effectively or outperforms analysts and other methods. This means that when designing AI solutions for AML, it is important to consult with the analysts, because their feedback can represent the only available source for understanding the system's performance.

Ensuring the security and ownership of data is also challenging. For machine learning applications, the use of public data resources is a common practice, and live data streams are also used. Relevant data are constantly required to enhance the performance of AML systems. However, financial institutions are always cautious about their data. Therefore, the quantity of open data in the AML field is limited. Another important problem is the lack of shared resources; apart from watch-lists and certain regulatory recommendations, institutions around the world do not maintain a shared data pool that can be used for the benefit of AML research. In such situations, the complexity of implementing GDPR's data privacy and ownership regulations may diminish the progress of AML research. Link analysis, sentiment analysis, and many other NLP- and knowledge-based

techniques are often used to reduce the high false positive rates in AML. The implementation of such components heavily depends upon public data. The implementation of data privacy and ownership policies can reduce the number of resources available for use. Furthermore, in terms of the explanations of decision-making processes, extreme care should be taken if the explanation is generated from a private data source. The possibility of disclosing important information is ever-present. Certain GDPR policies further require a "Data Protection Impact Assessment" to be adopted by organizations that process personal data, to identify and mitigate data-protection risks. This can be challenging when managing large datasets, given the complexity and unexpected uses of personal data.

The effectiveness of modern systems is somewhat limited by their ability to explain a specific decision that has been taken/predicted. Recent focus has explored the explainability of algorithms (Gunning 2017; Ehsan et al. 2018); however, the nature of explanation varies according to differences in the data and algorithms, and no common or standard framework of explanation has been implemented. For example, when operating image and convolutional neural networks (Kumar et al. 2017), explanations for a prediction can be presented in terms of low-level (e.g., edge, curve) and high-level (e.g., pattern) features obtained during learning. Interpreting the prediction is relatively easy for some (e.g., linear) models and hard for others [e.g., deep learning models such as long short-term memory (LSTM)]. A trade-off always exists between explanatory nature and model complexity.

Moreover, the scope of macro (i.e., the overall explanation for a decision) and micro (i.e., the explanation of machine learning components in the pipeline) explanations and their integration is another architectural challenge for agent-based AML solutions. Although frameworks can be easily constructed for rule-based systems, they also pose several challenges. For example, the explanations of complex rules for transaction monitoring can be subject to data privacy and sensitivity regulations. An explanatory framework for the data is also necessary. The legal and functional descriptions of the data must be transparent and visible. Legal descriptions refer to the data source and ownership, and functional descriptions refer to the data characteristics (e.g., whether the data contain any discriminatory properties, or the over- and/or under-representation of a certain group in a way that might harm real-world decisions). Machine learning—especially supervised machine learning—is completely data dependent, and a model can reflect inconsistencies/biases present in the data. A data-explanatory framework that responds to security, compatibility, and bias can play a vital role.

Some policies are readily translatable into technical requirements, and some are not. Explanations of a policy can be implemented in various stages of a system. However, when topics (e.g., "responsible AI") are discussed, it becomes unclear how and where such policies will be incorporated within a model/solution. Such philosophical debates give rise to many situations for which the correct answer is yet to found.

### 3.9 Survey papers

There are a number of literature reviews we would like to point out to the interested reader; these include Sudjianto et al. (2010), Rohit and Patel (2015), and Sharma and Panigrahi (2012). Similar research domains—such as fraud and auditing—also pertain to AML, and surveys of these domains can often be relevant; for example, Ngai et al. (2011), Pramanik et al. (2017), and Chintalapati and Jyotsna (2013).

## 4 An extension: a next-generation AML framework with AI

In this section, we propose a next-generation AML system using a multichannel fusion of state-of-the-art deep learning and NLP techniques. This novel AML framework applies and visualizes deep learning-driven NLP approaches in a distributed and scalable way, to enhance money laundering monitoring and investigation.

### 4.1 Applying deep learning and NLP to AML

In contrast to conventional machine learning approaches, deep learning methods can learn feature representations from raw data. In deep learning techniques, multiple layers of representation are learned from a raw data input layer, by using nonlinear manipulations on each representation learning level. Deep learning has outperformed many conventional machine learning approaches on designed features in various AI tasks, including natural language understanding, image recognition, and speech recognition (LeCun et al. 2015). NLP is a sub-domain of AI and frequently involves natural language understanding and generation; more specifically, it employs a set of techniques for syntax, semantics, and discourse analysis, text mining and classification, information extraction, and machine translation.

NLP and deep learning are already in use across many levels of AML regulatory compliance. However, their application is limited to decade-old techniques, and they fail to match the current pace of research. Deep learning frameworks have yet to be widely deployed in the AML community. In the recent literature, only one paper (Paula et al. 2016) has reported on the use of auto-encoders to detect suspicious transactions. Given that a vast amount of data is available across different financial institutions, deep learning techniques can prove useful.

#### 4.1.1 Entity recognition, entity resolution, and relation extraction

*Entity Recognition* is a set of algorithms capable of recognizing relevant entities (e.g., persons, positions, and companies) mentioned in an input text string. Relation extraction detects the relationship between two named entity nouns ($e1$, $e2$) in a given sentence—typically expressed as a triplet [$e1$, $r$, $e2$] (where, $r$ is a relationship between $e1$ and $e2$). Entity resolution (ER) determines whether references to

**Table 1** Example of the aliases and multiple dates of births registered by a single person

| Name | DOB | A.K.A |
|---|---|---|
| AL-TRABELSI | 20/05/1969 | ADEL Sassi, ADNAN Ben |
| MOURAD | 02/09/1966 | Salah, ADNAN Salah, AMOUR Bentaib |
| BEN ALI | 02/09/1964 | AROURI Farid, Ben ANAN Salah, BEN TAIEB, Arouri, Taoufik |
| BEN | 02/02/1963 | BRAHIM Aboue Chiba, BRAHIM Abouechiba, FAISEL Arouri |
| AL-BASHEER | 04/02/1965 | KAMEL Salam, MAGID Ben, MELLIT Hasnaoui, MELLIT Hasnaui |
| | 02/03/1965 | SALAH Adnan, SSASSI Maci and TAOUFIK, Arouri |

Source: https://www.gov.uk/government/publications/financial-sanctions-consolidated-list-of-targets/consolidated-list-of-targets

mentioned entities in various records and documents refer to the same or different entities. For example, the same person can be mentioned in different ways, and an organization could have different addresses.

Identifying inferences across different networks and semantic relationships between named entities is even more challenging when the amount of data grows. ER can reduce the task complexity by assigning canonicalized references to particular entities, or by de-duplicating and linking entities. The complexity of a network could be significantly reduced by de-duplication. For example, a seventh-order graph could be reduced to a much smaller three-order graph. Meanwhile, NLP and machine learning are employed in ER, as many other challenges appear in this procedure, including disambiguating confusion language, identifying abbreviations and truncation, recognizing various formatting, and spotting missing values. ER application to big data is more difficult since heterogeneity and cross-domain resolution is indispensable. Hence, ER techniques should be functionable with respect to big data techniques and distributed databases.

Generally, name screening is performed by comparing an aggregated list of suspected names, dates of birth, and other information; this is gathered both internally and externally using various data processes. ER can considerably assist in this area. Open ER remains a challenge, and very little work has been conducted towards its realization. Considering that suspicious entities hide their identities through several aliases (see Table 1), it is an open challenge to resolve and identify them.

#### 4.1.2 Sentiment analysis

Implementing sentiment analysis can be useful for AML; its primary role is to shorten the investigation period of a compliance officer. It can be applied at different levels, including the backlog management, client onboarding, and client profile-monitoring stages. The goal of a sentiment analysis system in this context is to monitor the sentiment trends associated with a client, to identify important patterns. When AML investigators identify a company that has potentially been involved in a suspicious transaction, they generally consult the Internet for evidence. Analyzing sentiment levels from news articles concerning a specific organization can reveal a great deal of evidence; for example, the consistent negative trends and

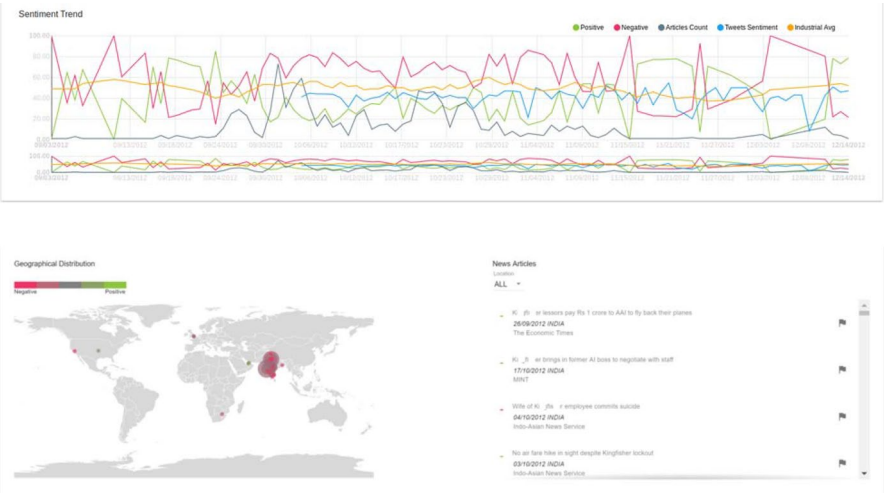


**Fig. 1** Sentiment Analysis for Kingfisher

negative keywords associated with the company Kingfisher may help human investigators quickly uncover the fraudulent or money laundering activities of the company (Fig. 1). AI-NLP based sentiment analysis can screen thousands of articles in seconds, significantly improving the investigation process in terms of efficiency and accuracy. In the following paragraph, we define a use case to show how and where a sentiment analysis model can enhance the performance of an AML system. Sentiment analysis can also be used in the client profile-monitoring and the client onboarding processes, to research and identify specific pain points of a client and their associations with negative articles.

Figure 1 depicts the sentiment analysis results of public domain news and reviews relating to Kingfisher Airlines. The data source for this study was a broad news corpus and a Twitter corpus. A recurrent neural network (RNN) was used to classify the news articles as positive, negative, or neutral. It also generated scores for each of the categories, between "positive" and "negative". The figure describes the overall sentiment trends present in the news pertaining to Kingfisher Airlines over the period from September 2012 to December 2012. The gray line in the graph indicates the number of articles/news per day, and the red line denotes the aggregated score for negative sentiment (per day). The size of the circles overlaid on the world-map represents the number of occurrences of the term "Kingfisher Airlines" in the news in that region, and the color conveys the sentiment. Green is positive, grey is neutral, and red is negative. A consistently high negative sentiment toward the organization over a certain period could indicate a potential candidate for blocking. Furthermore, applying aspect term extraction during this period can reveal the causes of the negative sentiment.

Sentiment analysis is a typical topic in NLP, and it has been employed in many different areas (Pang and Lee 2008). In terms of AI, numerous different techniques have been used for sentiment analysis, including SVM, conditional random fields,

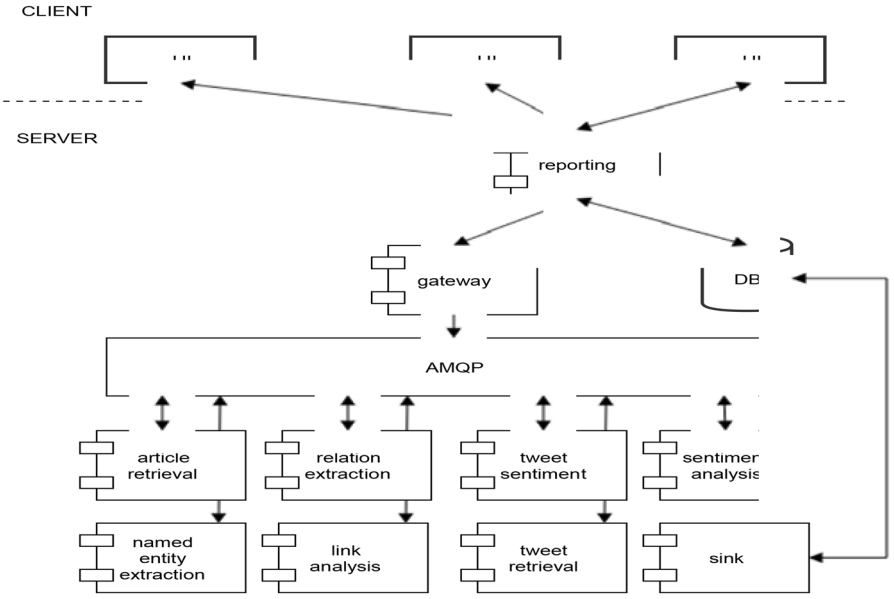


**Fig. 2** Explainable AI, human-in-the-loop technology, and NLP will be critical elements of future AML systems, addressing a number of challenges arising from ever-changing policies

and deep neural networks such as CNNs and RNNs. Significant accuracies in this task have been achieved (i.e., exceeding 80–90%) for classifying documents conveying positive and negative sentiments.

### 4.2 Overall pipeline of our system

Our proposed distributed AML framework integrates a set of NLP semantics tasks, including sentiment analysis, entity recognition, relation extraction, and entity linking analysis based on different dataset sources (e.g., news and Twitter posts), to provide additional references with which a human investigator can make a final decision. Each NLP module is evaluated on a task-specific dataset, and the overall experiments are performed on synthetic and real-world datasets. Feedback from AML practitioners suggests that our system can reduce the time and cost of operations by approximately 30% compared to their previous manual approaches toward AML investigation. To our best knowledge, the utilization of unstructured social and news data to facilitate AML has not been investigated in the recent literature.

#### 4.2.1 System architecture

The novel AML framework we propose is illustrated in Fig. 2. A distributed architecture integrating different NLP modules is adopted. The user interface (UI) displays the results from report services and passes the client's instructions to the

gateway layer via a reporting layer, thus triggering the processing pipeline. This reporting service allows the AML domain expert to retrieve supplementary social information relating to a specific transaction. The reported results can help a human analyst decide whether to approve or block the transaction. The information processing pipeline is a system of several different micro-services using different routings. The routing information is embedded in the messages that are exchanged, to make each component aware of the destination of the next message.

Banking and open data sources are handled in the system, and the datasets are governed, collected, managed, and stored in the database layer (DB). Banking data includes a wide variety of financial data, such as bank statements, know your customer (KYC), and client profiles. Open data refers to a collection of relevant financial news, social media posts, public financial reports, open-source fraud datasets, and so on.

#### 4.2.2 NLP and deep learning modules

*Relation Extraction*: Named entities are initially recognized and relationships are extracted therefrom. Seven types of named entities are defined in our system; namely person, organization, location, date, time, money, and miscellaneous. We implemented and enriched an attentive RNN framework (Lin et al. 2016) with both word-level and sentence-level attention layers in a relation-prediction learning procedure. We evaluated our model on the publicly available New York Times dataset, achieving an accuracy of 88.00% in terms of the P@100 measure.

Named entity recognition (NER) is performed using a combined strategy. We apply both the Stanford NER Recognizer and a neural NER [implemented using an LSTM-conditional random field (CRF) framework (Lample et al. 2016)]; then, we select out the specific types of entities we require from the recognized named entity combinations of the two models. The Stanford NER recognizer is a ready-to-use popular NLP toolkit. The LSTM-CRF model is an advanced deep learning-based method. It relies upon two sources of text information: character-based word representations (learned from the supervised corpus) and unsupervised word representations (learned from unannotated corpora). The LSTM-CRF-based models can achieve high NER prediction rates in four languages, without resorting to language-specific knowledge. The LSTM-CRF achieved best results (F1 = 90.94%) in English, German, Dutch, and Spanish NER (CoNLL-2002) tasks, compared with the models of Chiu and Nichols (2015) and Luo et al. (2015). A detailed comparison of LSTM-CRF and other models can be found in (Lample et al. 2016).

*Sentiment analysis* We used two different sentiment analysis (SA) models for our AML framework; namely, document-level and sentence-level models. The document-level model is a multi-channel CNN-based sentiment classifier (Kim 2014) that processes financial news articles. The sentence-level model is also a CNN-based classifier (Tang et al. 2014; Deriu et al. 2016); it is used for social media data. The adoption of CNN-based models was decided as a result of a comparison between them and bidirectional LSTM models; that is, both achieved similar performances,

thus we selected the CNN-based model for its superior model expandability. However, labeled financial resources are hard to obtain; to overcome this challenge and obtain a "goodness" prediction, we propose a voting scheme to annotate the training data; it is comprised of the following stages:

- Gathering financial news using a list of keywords[1] and news-searching Application Programming Interfaces (APIs).
- Passing the collected datasets (containing the headlines and first paragraphs of news articles) through publicly available sentiment APIs, to generate a sentiment score.
- Voting (via a voting mechanism) to obtain the final result in terms of positive or negative sentiments for each document.

Our document-level SA classifier was trained on 12,467 automatically labeled financial news articles, and it achieved an accuracy rate of 76.96%, which is comparable to a public sentiment API[2] based on the RT-polarity[3] data set. Our Twitter SA classifier was trained and evaluated on the SemEval-2016 task 4[4] dataset; it achieved an accuracy of 63.10%, comparable to the best system (63.30%) in the shared task (Han et al. 2018). In contrast to previous shared tasks, the SemEval-2016 task 4 is designed to predict the percentages of positive and negative tweets in given collections of tweets about a topic. Since many tweets may be related to a suspicious entity (e.g., an organization who may be involved in money laundering activities) in one period, this dataset is very useful for us to verify our SA models for AML scenarios.

#### 4.2.3 Evaluation and feedback from AML practitioners

As introduced in previous sections, we applied several validation and evaluation methods to different NLP models, where the Twitter SA model, news SA model, or attentive RE model achieve comparable results to the state-of-the-art in terms of accuracy. Recently, the entire AML system is being piloted and tested with our industry partners' global banking clients. It is being evaluated by professional AML practitioners for KYC investigations. Feedback from end users is that they are optimistic about reducing their time spent investigating red-alerted transactions by an average of 30%. We have been invited to deliver keynote talks concerning different aspects (not the entirety) of this system at highly respected events, including the Europe Financial Information Management Conference 2017, World Mobile Conference 2018, and others. In addition, our NLP models have also been utilized by clients in different domains; for instance, our SA models were employed for trend monitoring of brand reputation (Han et al. 2018).

[1] https://www3.nd.edu/~mcdonald/Word_Lists.html.

[2] We use https://www.ibm.com/watson/alchemy-api.html and it achieves 75.56% in terms of accuracy.

[3] https://www.cs.cornell.edu/people/pabo/movie-review-data/rt-polaritydata.README.1.0.txt.

[4] Prediction of five-point scale polarity of a tweet.

### 4.3 How our approach addresses policy limitations

#### 4.3.1 Data governance and ownership with blockchain

During the onboarding process, financial institutions will typically collect, verify, and screen customer identities, to comply with AML and KYC regulations. Blockchain can validate and verify entities within a system through the use of digital certificates, which can be associated with data stored on the blockchain and linked to official or government ID.

Once transaction data has been added and verified on the blockchain, it is immutable; that is, it cannot be removed, only updated. This creates a clear and auditable record of transactions, facilitating the tracking of transaction sources and destinations within the system, thereby making it more difficult to launder money. This also safeguards against double-spending and transaction reversals.

Consortium blockchains have a number of applications and possibilities. For example, through them it would be possible to create a partially decentralized, permissioned blockchain in which a number of international financial institutions are stakeholders, requiring a consensus vote between them to validate transactions; this would simultaneously allow direct oversight by a number of third parties (government bodies, international regulators, concerned citizens, etc.)

The increased transaction visibility offered by blockchain facilitates algorithmic approaches to identifying patterns and high-level monitoring, allowing regulators to determine risks more effectively.

The advent of blockchain technologies—in particular, due to the attention garnered by Bitcoin—has led to a rapid growth in FinTech startups. Several of these are developing new blockchain architectures specifically designed for payment transactions (Bitcoin, Ripple) or as general-purpose ledgers (Ethereum, Hyperledger). Others focus on providing tools for forensic analysis, fraud prevention, and regulatory compliance (AML, KYC) on existing blockchains such as Bitcoin and Ethereum; this category includes Elliptic, Chainalysis, Coinfirm, Scorechain, and IdentityMind.

A blockchain is a database in which copies of the data are shared across multiple locations (nodes) and validated without a central authority. It is capable of storing any type of data, although its "write-once, read-only" design has advantages in functioning as a ledger for storing financial data with strong security and transparency requirements (e.g., transaction records, customer identities, company shares, title deeds, etc.). Data is stored as "blocks," and each block is cryptographically validated by other nodes in the system, which must reach consensus before the block is added to the chain of previous transactions. This validated blockchain contains a copy of all previous transactions executed in the system and is visible to all participating entities in the blockchain.

Blockchain implementation is flexible, making it suitable for a number of applications. Besides the underlying technical details (e.g., programming language, cryptography, blocksize), the two primary implementation design choices are whether it is a public or private (or consortium) blockchain, and whether it is permissioned or permission-less. These choices reflect the application requirements for anonymity, efficiency, and transparency; however, each also entails certain trade-offs.

In a public blockchain, anybody with a computer and an internet connection can connect to the network, set up a node, and obtain the complete transaction history of that blockchain. This generally produces a high level of data redundancy and a lower efficiency. As a corollary, the high visibility makes it extremely secure and tamper-proof, although with disadvantages in terms of data privacy. By contrast, a private blockchain restricts write-access to a primary entity which validates and writes each block. This increases the efficiency and speed of transactions because they do not require validation from multiple nodes, though sacrifices are made in terms of the visibility and security of the system. However, read-access can still be extended to third parties, allowing for procedures such as auditing. The third option is a partially private or consortium blockchain, in which a group of entities (i.e., stakeholders, typically pre-selected) share access, with none possessing full control. This design has clear applications for financial institutions.

The second design choice is permissioned versus permission-less. A permissioned blockchain adds a cryptographic access control layer, so that only those users with the proper permissions can validate transactions and participate in the consensus vote. In a permission-less blockchain, anyone on the network can add a transaction—though it still requires validation—and participate in the consensus vote.

Blockchain technology creates secure distributed ledgers that can store transaction records, customer identities, and other forms of data. It addresses many of the problems currently associated with AML compliance, particularly those relating to data protection and governance; furthermore, it can improve upon current practices in areas such as client onboarding, identity verification, transaction monitoring, and reporting.

The openness and transparency of blockchain transactions is particularly interesting. Recently, an analysis conducted by Griffin and Shams (2018) investigated whether Tether—said to be pegged to the US dollar—had an influence on the Bitcoin price. The crucial point is the algorithms and methods used to decipher money flows between major cryptocurrency exchanges. These determine seed wallet addresses for a number of exchanges and extrapolate other associated wallets. Then, they monitor the flow of coins to and from addresses at critical time points. This information is available to all who search for it, owing to the public blockchain record of transaction details. This would greatly benefit AML practices; indeed, it is believed "to be under consideration in many governmental and law enforcement agencies. Advances have been made in cryptocurrency research, aiming to thwart the traceability of transactions from sender to receiver; these include the coins Monero, which utilize ring signatures, and ZCash, which uses a privacy focused protocol referred to as sk-SNARKS.

### 4.3.2 Explainable decisions and humans-in-the-loop

Because communication with an analyst are critical, it is also important to consider the way in which information is presented to the analysts for evaluation. We present the user interface we have created as an example of how to gather and present relevant information to an analyst. We explain how we implement NLP methods (including entity extraction and sentiment analysis) to augment the analysts'

available information concerning a transaction, without requiring them to conduct the research.

Communication with analysts is of utmost importance when designing any AML system, because the users make the final decision. Explainable AI methods operate by providing the user with clear information on why a prediction was made (e.g., why the system believes a transaction is suspicious), to assist the user in making a decision and further the user's understanding of the technology. It is important that any system is able to explain its decisions in a user-friendly way. We have built a system to communicate with a user; it compiles the relevant information for the review of suspicious transactions using NLP techniques, including ER and relation extraction. For any flagged transaction, our system compiles the most relevant news articles pertaining to the company or individuals involved in the transaction. It summarizes the sentiments of these articles and provides a link analysis between the transaction and known criminal entities.

European policy emphasizes the need for financial institutions to provide explainable and human-authorized decisions. It is critical that any future AML method incorporates a human analyst and ensures that the analyst clearly understands the information they are being presented. A black-box system that labels a transaction as "fraudulent" with no further insight is unacceptable. In this paper, we propose a system structure designed around the analyst. Our human-in-the-loop system learns from the decisions of analysts and the feedback they give, to minimize the number of transactions the analyst must screen for fraud detection. Unlike the current rule-based methods, our system continuously learns to identify frauds, reducing the number of transactions an analyst must review over time. This continuous learning will also make the system adaptable to changes in fraud and policy.

Furthermore, in the case of a transaction-monitoring system, when an algorithm predicts an innocent transaction as malicious using metrics that resemble the historical data, it becomes unclear whether the boundary between prediction and active avoidance could be used to achieve a high true positive rate. To manage such situations, certain ethical standards must be followed when designing AI solutions to a specific problem. In terms of ethics, it is inevitable that some human elements must be present in the system design, to evaluate and correct the decisions of the predictive model. Inconsistencies—such as those in the data or (perhaps) a prediction bias due to discriminatory properties of the data—can only be addressed by human evaluators. For example, when considering whether a transaction is fraudulent or not, an AML system can present its predictions with an explanation to the end user; a final decision is made by the human (who may or may not support the system prediction), eventually this decision is back-propagated to the system to improve its decision-making capabilities. A similar paradigm can also be adopted for name-screening solutions.

In this context, the importance of human–computer interaction (HCI) plays a vital role. AML compliance depends upon gathering evidence against suspicious transactions. In practice, the process is laborious and complicated. For example, it involves the use of search engines—to gather and filter data manually in the operational layer—and the investigation of similar transactions from a historical repository. Moreover, the manner in which interactions occur between a compliance officer and

the system is efficient. The incorporation of HCI can bring significant improvements in this area.

We argue that future AML systems should not be linear; they should be cyclical systems in which the automated methods communicate with and learn from the analysts. Due to the technical challenges (caused by the lack of labeled transactions as well as new policies that require a human-generated explanation for transactions that are denied and accounts that are shut down), it is critical that the analyst and automatic method work together.

Explainability is another issue that must be addressed when managing GDPR regulations. AI systems implemented at different levels of current AML solutions (e.g., outlier-detection algorithms in transaction monitoring systems) can present their predictions via a visual interface as an explanation (Sect. 3.5). For name-screening solutions, explanations can be generated by displaying the similarity metrics against different screening criteria. However, the means of describing their prediction will depend on the task and the algorithm.

## 5 Conclusion

In this paper, we review the state-of-the-art AI methods for AML and extend the discussion by proposing a framework for next-generation AML; this framework utilizes advanced NLP and deep learning techniques. Numerous researchers have examined ways to increase the true fraud prediction rate whilst reducing the false suspicious transaction alarm rate. We develop a framework that utilizes unstructured external information to assist domain experts, aiming to decrease the workload for the human investigator.

Our solution extracts information from multiple social networks, which are relevant in many money laundering sectors. It would function well even if a money launder were to realize that their social networks were under inspection, because their networks and behaviors cannot be easily changed.

Researchers have found that (traditional) social network analysis is an effective and important element of AML systems. Our new framework strengthens the analysis by introducing social media and other web sources to the detection process of money laundering activities. In addition, our system is not intended to replace the current AML systems. It is complimentary to the current solutions; it augments the existing systems by providing auxiliary information in a clear, concise, and consumable format. The new framework may be used alongside a bank's sensitive AML alarm system, providing human experts with additional assistance for decision-making following the detection of a suspicious transaction. Meanwhile, graph learning and NLP techniques are recent advances that could become standard approaches for AML. Our next-generation system can boost the efficiency of money laundering detection without significant new capital investment from financial institutions, effectively reducing their potential reputational risk and cost.

Accessing AML data sets is an existing unsolved problem for the AML research community. There are a limited number of annotated money laundering datasets publicly available; this is a major problem that holds back AML research because

a lot of AI approaches cannot be directly applied to transaction data and must rely on these datasets. This is especially true for deep learning approaches—large, real-world annotated datasets are essential to their application. Certain deep learning approaches (such as reinforcement learning) have achieved competitive performances in other domains, but they are yet to be tested for AML. Meanwhile, unsupervised machine learning approaches have been overlooked in the literature. Anonymized public AML datasets (that conceal clients' information) would benefit the whole AML research community. Ideally, these would be sourced from large banks. Depending on the data types, further anonymization may be needed to prevent the disclosure of clients' identities and other information.

**Acknowledgements** This work was partially supported by Enterprise Ireland (grant number IP20170626). We would like to thank Accenture Applied Intelligence, Fraud and Risk Analytics and Technology Labs teams for inspiring conversations and support.

## References

Barnett, M. L. (2014). Why stakeholders ignore firm misconduct: a cognitive view. *Journal of Management, 40*(3), 676–702.

Bavelas, A. (1950). Communication patterns in task-oriented groups. *The Journal of the Acoustical Society of America, 22*(6), 725–730.

Bernile, G., & Gregg, A. J. (2009). The impact of the options backdating scandal on shareholders. *Journal of Accounting and Economics, 47*(1–2), 2–26.

Chintalapati, S. S., & Jyotsna, G. (2013). Application of data mining techniques for financial accounting fraud detection scheme. *International Journal of Advanced Research in Computer Science and Software Engineering, 3*, 717–724.

Chiu, J. P. C., & Nichols, E. (2016). Named entity recognition with bidirectional LSTM-CNNs. *Transactions of the Association for Computational Linguistics, 4*, 357–370.

Colladon, A. F., & Remondi, E. (2017a). Using social network analysis to prevent money laundering. *Expert Systems with Applications, 67*, 49–58. https://doi.org/10.1016/j.eswa.2016.09.029.

Colladon, A. F., & Remondi, E. (2017b). Using social network analysis to prevent money laundering. *International Journal of Expert Systems with Applications, 67*, 49–58.

Deriu, J., Gonzenbach M, Uzdilli F, Lucchi A, De Luca V, and Jaggi M. 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*, 1124–8. EPFL-CONF-229234.

Drezewski, R., Sepielak, J., & Filipkowski, W. (2015). The application of social network analysis algorithms in a system supporting money laundering detection. *Information Scienece, 295*, 18–32.

Ehsan, U., Harrison, B., Chan, L. & Riedl, M.O. (2018). Rationalization: a neural machine translation approach to generating natural language explanations. *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society (AIES'18)*, New Orleans, LA, USA, 2–3 February 2018, 81–87.

Gallo, P. A., & Juckes, C. C. (2005). Threshold transaction disclosures: access on demand through latent disclosure rather than reporting. *Journal of Money Laundering Control, 8*, 328–334.

Gao, S., & Xu, D. (2007). Conceptual modeling and development of an intelligent agent-assisted decision support system for anti-money laundering. *Expert Systems with Applications, 36*, 1493–1504.

Gao, Z., & Ye, M. (2007). A framework for data mining-based anti-money laundering research. *Journal of Money Laundering Control*, *10*(2), 170–179.

Gao, S., & Xu, D. (2010). Real-Time Exception Management Decision Model (RTEMDM): Applications in Intelligent Agent-Assisted Decision Support in Logistics and Anti-Money Laundering Domains. *Proceedings of the 43rd Hawaii International Conference on System Sciences*, Honolulu, HI, USA, 2010, 1–10.

Gao, S., Xu, D., Wang, H., & Wang, Y. (2006). Intelligent Anti-money Laundering System. *Proceedings of the IEEE International Conference on Service Operations and Logistics and Informatics*, Shanghai, Peoples Republic of China, 21–23 June 2006, 851–856.

Goldberg, H.G., & Wong, R.W.H. (1998). Restructuring transactional data for link analysis in the FinCEN AI System. *AAAI Technical Report*, *FS-98-01*, 38–46.

Gombiro, C., & Jantjies, M. (2015). A conceptual framework for detecting financial crime in mobile money transactions. *Journal of Governance and Regulation*, *4*(4, continued 6),727–734.

Goodman, B., & Flaxman, S. (2016). European union regulations on algorithmic decision-making and a 'right to explanation. arXiv:1606.08813

Griffin, J.M, & Shams, A. (2018). Is Bitcoin Really Un-Tethered? Working paper. Available at SSRN: https://ssrn.com/abstract=3195066.

Gunning, D. (2017). Explainable artificial intelligence (Xai), DARPA/I2O. Retrieved from https://www.darpa.mil/attachments/XAIProgramUpdate.pdf.

Han, J., Barman, U., Hayes, J., Du, J., Burgin, E., Wan, D. (2018). NextGen AML: distributed deep learning based language technologies to augment anti money laundering investigation. *Proceedings of the ACL 2018, System Demonstrations,* Melbourne, Australia, July 2018, 37–42.

Hand, D.J., & Weston, D.J. (2008). *Statistical techniques for fraud detection, prevention, and assessment*. In F. Fogelman-Soulie, D. Perrotta, J. Piskorski, & R. Steinberger (Eds.) Mining massive data sets for security (pp. 257–270), Amsterdam: IOS Press.

Hanneman, R.A, & Riddle, M. (2005). *Introduction to social network methods.* Riverside, CA: University of California, Riverside.

Helmy, T. H. E., Abd-ElMegied, M. Z., Sobh, T. S., & Badran, K. M. S. (2014). Design of a monitor for detecting money laundering and terrorist financing. *International Journal of Computer Networks and Applications*, *1*(1), 15–25.

Hendrickx, I., Kim, S.N., Kozareva, Z., Nakov, P., Séaghdha, D.O., Padó, S., Pennacchiotti, M., Romano, L., & Szpakowicz, S. (2010). Semeval-2010 Task 8: multi-way classification of semantic relations between pairs of nominals. *Proceedings of the 5th International Workshop on Semantic Evaluation, ACL 2010*, Uppsala, Sweden, 15–16 July 2010, 33–38.

Huff, R. M. (2014). *Money laundering. The Encyclopedia of Criminology and Criminal Justice*. New York: Wiley.

Irish Examiner. (2016). Retrieved from https://www.irishexaminer.com/breakingnews/ireland/ulster-bank-hit-with-33mfine-for-anti-money-laundering-weaknesses-761956.html.

Johnson, W. C., Xie, W., & Yi, S. (2014). Corporate fraud and the value of reputations in the product market. *Journal of Corporate Finance, 25*, 16–39.

Jorisch, A. (2009). *Tainted Money: Are We Losing the War on Money Laundering and Terrorism Financing?* Arlington: Red Cell Intelligence Group.

Kamarinou, D., Millard, C., Singh, J. (2016). Machine learning with personal data: Profiling, decisions and the EU General Data Protection Regulation. Working paper. Available at http://www.mlandthelaw.org/papers/kamarinou.pdf.

Kannan, S., & Somasundaram, K. (2017). Autoregressive-based outlier algorithm to detect money laundering activities. *Journal of Money Laundering Control*, *20*(2), 190–202.

Karpoff, J. M., Lee, D. S., & Martin, G. S. (2008). The cost to firms of cooking the books. *Journal of Financial and Quantitative Analysis, 43*(3), 581–611.

Khan, N., & Haider, S. (2013). A Bayesian approach for suspicious financial activity reporting. *International Journal of Computers and Applications, 35*, 181–187.

Kim, Y. (2014). Convolutional neural networks for sentence classification. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, 25–29 October 2014, 1746–1751.

Kingdon, J. (2004). AI fights money laundering. *IEEE Intelligent Systems, 19*, 87–89.

Kolhatkar, J.S., Fatnani, S.S., Yao, Y., & Matsumoto, K. (2014). Multi-channel data driven, real-time anti-money laundering system for electronic payment cards. U.S. Patent No. US 8,751,399 B2. Available at: https://patentimages.storage.googleapis.com/20/52/22/4f12c57929b368/US8751399.pdf.

Kumar, D., Wong, A., & Taylor, G.W. (2017). Explaining the unexplained: a class-enhanced attentive response (CLEAR) approach to understanding deep neural networks. https://arxiv.org/abs/1704.04133.

Kuner, C., Svantesson, D. J. B., Cate, F. H., Lynskey, O., & Millard, C. (2017). Machine learning with personal data: is data protection law smart enough to meet the challenge? *International Data Privacy Law, 7*(1), 1–2.

Lai, K. (2018). Blockchain as AML tool: a work in progress. *International Financial Law Review.*

Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., & Dyer, C. (2016). Neural architectures for named entity recognition. https://arxiv.org/abs/1603.01360.

Larik, A. S., & Haider, S. (2010). Clustering based anomalous transaction reporting. *Procedia Computer Science, 3*, 606–610.

Le-Khac, N.-A., Kechadi, T. (2010). Application of data mining for anti-money laundering detection: A case study. *Proceedings of 2010 IEEE International Conference on Data Mining Workshops*, Sydney, Australia, December 2010, 577–584.

Le-Khac, N.-A., Markos, S., Kechadi, M.-T. (2009). Towards a new data mining-based approach for anti-money laundering in an international investment bank. In *International Conference on Digital Forensics and Cyber Crime*.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436.

Lin, Y., Shen, S., Liu, Z., Luan, H., Sun, M. (2016). Neural relation extraction with selective attention over instances. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Berlin, Germany, August 2016, 2125–2133.

Liu, X., Zhang, P., & Zeng, D. (2008). Sequence matching for suspicious activity detection in anti-money laundering. In *Intelligence and Security Informatics Workshops*.

Lopez-Rojas, E.A., & Axelsson, S. (2012a). Money laundering detection using synthetic data. *Proceedings of the 27th Annual Workshop of the Swedish Artificial Intelligence Society (SAIS)*, Örebro, Sweden, 14–15 May 2012, 33–40.

Lopez-Rojas, E.A., & Axelsson, S. (2012b). Multi agent based simulation (MABS) of financial transactions for anti money laundering (AML). *Proceedings of the 17th Nordic Conference on Secure IT System*, Karlskrona, Sweden, October 31 – November 2, 2012, 25–32.

Lucia, D.P., Donato, M. (2009). The risk-based approach in the new european anti-money laundering legislation: a law and economics view. *Review of Law & Economics*, 5(2), 931–52.

Luo, G., Huang, X., Lin, C.Y., & Nie, Z. (2015). Joint entity recognition and disambiguation. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Lisbon, Portugal, September 2015, 879–888.

Lv, L.-T., Ji, N., & Zhang, J.L. (2008). A RBF neural network model for anti-money laundering. *Proceedings of the 2008 International Conference on Wavelet Analysis and Pattern Recognition*, Hong Kong, China, 30–31 August 2008, 209–215.

Martin, B. (2017, January 31). Deutsche Bank Hit with £500m Money Laundering Fines. *The Telegraph*. Retrieved from https://www.telegraph.co.uk/business/2017/01/31/deutsche-bank-hit-500m-money-laundering-fines/.

Mei, D., Ye, Y., & Gao, Z. (2014). Literature review of international anti-money laundering research: a scientometrical perspective. *Open Journal of Social Sciences, 2*, 111–120.

Moustafa, T. H., El-Megeid, M. Z. A., Sobh, T. S., & Shafea, K. M. (2015). Anti money laundering using a two-phase system. *Journal of Money Laundering Control*, *18*(3), 304–329.

Murphy, D. L., Shrieves, R. E., & Tibbs, S. L. (2009). Understanding the penalties associated with corporate misconduct: an empirical examination of earnings and risk. *Journal of Financial and Quantitative Analysis, 44*(1), 55–83.

Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The applications of data mining techniques in financial fraud detection: a classification framework and an academic review of literature. *Decision Support Systems, 50*, 559–569.

Owen, S., Anil, R., Dunning, T., & Friedman, E. (2011). *Mahout in Action*. Greenwich: Manning Publications Co.

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval*, *2*(1–2), 1–135.

Parkman, T. (2012). *Mastering Anti-Money Laundering and Counter-Terrorist Financing: A Compliance Guide for Practitioners*. Harlow, England: Pearson.

Paula, E.L, Ladeira, M., Carvalho, R.N., & Marzagão, T. (2016). Deep learning anomaly detection as support fraud investigation in Brazilian exports and anti-money laundering. *Proceedings of the 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*, Anaheim, CA, USA, 18–20 Dec. 2016, 954–960.

Pramanik, M. I., Lau, R. Y. K., Yue, W. T., Ye, Y., & Li, C. (2017). Big data analytics for security and criminal investigations". *WIREs Data Mining Knowledge Discovery, 7*, e1208.

Rajput, Q., Khan, N. S., Larik, A., & Haider, S. (2014). Ontology based expert-system for suspicious transactions detection". *Computer and Information Science, 7*, 103.

Raza, S., & Haider, S. (2010). Suspicious activity reporting using dynamic bayesian networks. *Procedia Computer Science, 3*, 987–991.

Rohit, K. D., & Patel, D. B. (2015). Review on detection of suspicious transaction in anti-money laundering using data mining framework. *Journal for Innovative Research in Science and Technology, 1*, 129–133.

Senator, T. E., Goldberg, H. G., Wooton, J., Cottini, M. A., Khan, A. F. U., Klinger, C. D., et al. (1995). The financial crimes enforcement network Ai System (Fais) identifying potential money laundering from reports of large cash transactions. *AI Magazine, 16*, 21.

Sharma, A., & Panigrahi, P. K. (2012). A review of financial accounting fraud detection based on data mining techniques. *International Journal of Computer Applications*, *39*(1), 37–47.

Sudjianto, A., Nair, S., Yuan, M., Zhang, A., Kern, D., & Cela-Diaz, F. (2010). Statistical methods for fighting financial crimes". *Technometrics, 52*, 5–19.

Tang, D., Wei, F., Yang, N., Zhou, M., Liu, T., Qin, B. (2014). Learning sentiment-specific word embedding for twitter sentiment classification. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, Baltimore, Maryland, USA, June 23–25 2014, 1555–1565.

Tang, J., & Yin, J. (2005). Developing an intelligent data discriminating system of anti-money laundering based on SVM. *Proceedings of the 2005 International Conference on Machine Learning and Cybernetics*, Guangzhou, China, 18–21 Aug. 2005, 3453–3457.

Thompson, M., & Perez, E. (2017, July 1). BNP Paribas to Pay Nearly $9 Billion Penalty. *CNN Business*. Retrieved from https://money.cnn.com/2014/06/30/investing/bnp-paribas-sanctions-fine/.

Titcomb, J. (2014, Aug 9). Standard chartered pays $300m over money laundering failures. *The Telegraph*. Retrieved from https://www.telegraph.co.uk/finance/newsbysector/banksandfinance/11044353/Standard-Chartered-pays-300m-over-money-laundering-failures.html.

Unger, B., & Van Waarden, F. (2009). How to dodge drowning in data? Rule- and risk-based anti money laundering policies compared. *Review of Law and Economics, 52*, 953–985.

Viswanatha, A., & Wolf, B. (201). HSBC to pay $1.9 Billion U.S. Fine in Money-Laundering Case. *Reuters*. Retrieved from https://uk.reuters.com/article/us-hsbc-probe-idUSBRE8BA05M20121211.

Wang, S.-N., & Yang, J.-G. (2007). A money laundering risk evaluation method based on decision tree. *Proceedings of 2007 International Conference on Machine Learning and Cybernetics*, Hong Kong, China, 19-22 August 2007, 283–286.

Weber, M., Chen, J., Suzumura, T., Pareja, A., Ma, T., Kanezashi, H., Kaler, T., Leiserson, C.E., & Schardl, T.B. (2018). Scalable graph learning for anti-money laundering: a first look. arXiv:1812.00076.

White, T. (2009). *Hadoop: The Definitive Guide* (1st ed.). Newton: O'Reilly Media Inc.

Yu, T., Sengul, M., & Lester, R. H. (2008). Misery loves company: the spread of negative impacts resulting from an organizational crisis. *Academy of Management Review, 33*(2), 452–472.

Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S., Stoica, I. (2010). Spark: cluster computing with working sets. *Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing (HotCloud'10)*, Boston, Massachusetts, USA, June 2010, 1–7.

Zengan, G. (2009). Application of cluster-based local outlier factor algorithm in anti-money laundering. *Proceedings of 2009 International Conference on Management and Service Science*, Wuhan, China, 20–22 September 2009, 1–4.

Zhang, Z., Yu, P.S., & Salerno, J.J. (2003). Applying data mining in investigating money laundering crimes. *Proceedings of the ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Washington, DC, USA, August 24–27, 2003, 747–752.