

TOPICAL REVIEW

Fighting Money Laundering With Statistics and Machine Learning

RASMUS INGEMANN TUFFVESON JENSEN^{1,2}
AND ALEXANDROS IOSIFIDIS¹, (Senior Member, IEEE)

¹DIGIT, Department of Electrical and Computer Engineering, Aarhus University, 8000 Aarhus, Denmark

²AML and Fraud, Spar Nord Bank, 9100 Aalborg, Denmark

Corresponding author: Rasmus Ingemann Tuffveson Jensen (rasmus.tuffveson.jensen@ece.au.dk)

ABSTRACT Money laundering is a profound global problem. Nonetheless, there is little scientific literature on statistical and machine learning methods for anti-money laundering. In this paper, we focus on anti-money laundering in banks and provide an introduction and review of the literature. We propose a unifying terminology with two central elements: (i) client risk profiling and (ii) suspicious behavior flagging. We find that client risk profiling is characterized by diagnostics, i.e., efforts to find and explain risk factors. On the other hand, suspicious behavior flagging is characterized by non-disclosed features and hand-crafted risk indices. Finally, we discuss directions for future research. One major challenge is the need for more public data sets. This may potentially be addressed by synthetic data generation. Other possible research directions include semi-supervised and deep learning, interpretability, and fairness of the results.

INDEX TERMS Anti-money laundering, know-your-client, machine learning, literature review.

I. INTRODUCTION

Officials from the United Nations Office on Drugs and Crime estimate that money laundering amounts to 2.1-4% of the world economy [1]. The illicit financial flows help criminals avoid prosecution and undermine public trust in financial institutions [2], [3], [4]. Multiple intergovernmental and private organizations assert that modern statistical and machine learning methods hold great promise to improve anti-money laundering (AML) operations [5], [6], [7], [8], [9]. The hope, among other things, is to identify new types of money laundering and allow a better prioritization of AML resources. The scientific literature on statistical and machine learning methods for AML, however, remains relatively small and fragmented [10], [11], [12].

The international framework for AML is based on recommendations by the Financial Action Task Force (FATF) [13]. Within the framework, any interaction with criminal proceeds practically corresponds to money laundering from a bank perspective (regardless of intent or transaction complexity) [14]. Furthermore, the framework requires that banks:

- 1) know the identity of, and money laundering risk associated with, clients, and
- 2) monitor and report suspicious behavior.

Note that we, to reflect FATF's recommendations, are intentionally vague about what constitutes "suspicious" behavior.

To comply with the first requirement, banks ask their clients about identity records and banking habits. This is known as know-your-customer (KYC) information and is used to construct risk profiles. The profiles are, in turn, often used to determine intervals for ongoing due diligence, i.e., checks on KYC information.

To comply with the second requirement, banks use electronic AML systems to raise alarms for human inquiry. Bank officers then dismiss or report the alarms to national financial intelligence units (i.e., authorities). The process is illustrated in Figure 1. Traditional AML systems rely on predefined and fixed rules [15], [16]. Although the rules are formulated by experts, they are essentially "if-this-then-that" statements; easy to interpret but inefficient. Indeed, over 98% of all AML alarms can be false positives [17]. Banks are not allowed to disclose information about alarms and generally receive little feedback on filled reports. Furthermore, money launderers may change their behavior in response to AML efforts. For

The associate editor coordinating the review of this manuscript and approving it for publication was Derek Abbott¹.




FIGURE 1. Process of an AML alarm. First, an AML system raises the alarm. A bank officer then reviews it. Finally, it is either dismissed or reported to authorities.

instance, banks in the United States must, by law, report all currency transactions over \$10,000 (regardless of whether they constitute money laundering or not) [18]. In response, money launderers may employ smurfing (i.e., split up large transactions). Finally, as money laundering has no direct victims, it can potentially go undetected for longer than other types of financial crime (e.g., credit card or wire fraud).

In this paper, we focus on AML in banks and aim to provide a technical review that researchers and industry practitioners (statisticians and machine learning engineers) can use as a guide to the current literature on statistical and machine learning methods for AML in banks. Furthermore, we aim to provide a terminology that can facilitate policy discussions, and to provide guidance on open challenges within the literature. To achieve our aims, we (i) propose a unified terminology for AML in banks, (ii) review selected exemplary methods, and (iii) present recent machine learning concepts that may improve AML.

The rest of the paper is organized as follows. Section II presents our terminology, distinguishing between (i) client risk profiling and (ii) suspicious behavior flagging. Section III then reviews the literature on client risk profiling, while Section IV reviews the literature on suspicious behavior flagging. Note that both Sections III and IV contain subsections that further distinguish between unsupervised and supervised methods. Next, Section V discusses future research directions. Finally, Section VI concludes the paper.

II. TERMINOLOGY

Inspired by FATF’s recommendations, we argue that banks face two principal data analysis problems in AML: (i) client risk profiling and (ii) suspicious behavior flagging. We use these to structure our terminology and review. A related topic, not discussed here, concerns how authorities treat AML reports (see, for instance, Savage et al. [19], Drezewski et al. [20], Li et al. [21], or Baltoi et al. [22]). We further make a distinction between unsupervised and supervised methods. Unsupervised methods utilize data sets on the form $\{\mathbf{x}_c \mid c = 1, \dots, n\}$ where n denotes some number of clients. Supervised methods, by contrast, utilize data

sets $\{(\mathbf{x}_c, y_c) \mid c = 1, \dots, n\}$ where some labels (e.g., risk scores) y_c are given.

A. CLIENT RISK PROFILING

Client risk profiling is used to assign general risk scores to clients. Let $\mathbf{x}_c \in \mathbb{R}^d$ be a vector of features specific to client c and \mathcal{P} be a generic set. A client risk profiling is a mapping

$$\rho : \mathbb{R}^d \rightarrow \mathcal{P}, \tag{1}$$

where $\rho(\mathbf{x}_c)$ captures the *money laundering risk* associated with client c . For example, we may have $\mathcal{P} = \{L, M, H\}$, where L symbolizes low risk, M symbolizes medium risk, and H symbolizes high risk. We stress that client risk profiling in our terminology is characterized by working on the client, not transaction, level.

B. SUSPICIOUS BEHAVIOR FLAGGING

Suspicious behavior flagging is used to raise alarms on clients, accounts, or transactions. Consider a setup where client c has $a = 1, \dots, A_c$ accounts. Furthermore, let each account (c, a) have $t = 1, \dots, T_{(c,a)}$ transactions and let $\mathbf{x}_{(c,a,t)} \in \mathbb{R}^d$ be some features specific to transaction (c, a, t) . An AML system is a function

$$s : \mathbb{R}^d \rightarrow \{0, 1\}, \tag{2}$$

where $s(\mathbf{x}_{(c,a,t)}) = 1$ indicates that an alarm is raised on transaction (c, a, t) . Multiple approaches may be used to construct an AML system. Regardless of approach, we argue that all AML systems are built on one fundamental premise. To cite Bolton and Hand [23]: “...given that it is too expensive to undertake a detailed investigation of all records, one concentrates investigation on those thought most likely to be fraudulent.” Thus, a good AML system needs to model the probability

$$F(\mathbf{x}_{(c,a,t)}) = P(y_{(c,a,t)} = 1 \mid \mathbf{x}_{(c,a,t)}), \tag{3}$$

where $y_{(c,a,t)} = 1$ indicates that transaction (c, a, t) should be reported for money laundering (with $y_{(c,a,t)} = 0$ otherwise). We may then raise alarms given some threshold value $\epsilon \geq 0$ and an indicator function $s(\mathbf{x}_{(c,a,t)}) = 1_{\{F(\mathbf{x}_{(c,a,t)}) \geq \epsilon\}}$.

It can be difficult to determine if a transaction, in itself, is money laundering. As a remedy, the level of analysis may be changed (see Figure 2). We may, for instance, consider account features $\mathbf{x}_{(c,a)} \in \mathbb{R}^d$ that summarize all activity on account (c, a) . Alternatively, we may consider the set of all feature vectors $\mathcal{X}_{(c,a)} = \{\mathbf{x}_{(c,a,1)}, \dots, \mathbf{x}_{(c,a,T_{(c,a)})}\}$ for transactions $t = 1, \dots, T_{(c,a)}$ made on account (c, a) . Defining $y_{(c,a)} \in \{0, 1\}$ in analogy to $y_{(c,a,t)}$, we may then model

$$F(\mathbf{x}_{(c,a)}) = P(y_{(c,a)} = 1 \mid \mathbf{x}_{(c,a)}) \tag{4}$$

or

$$F(\mathcal{X}_{(c,a)}) = P(y_{(c,a)} = 1 \mid \mathcal{X}_{(c,a)}), \tag{5}$$

i.e., the probability that account (c, a) should be reported for money laundering given $\mathbf{x}_{(c,a)}$ or $\mathcal{X}_{(c,a)}$. Similarly, we could




FIGURE 2. A single client may hold multiple accounts in a bank, each facilitating numerous transactions. When doing suspicious behavior flagging, alarms may be raised at the client, account, or transaction level (or a combination of them).

raise alarms directly at the client level, modeling

$$F(\mathbf{x}_c) = P(y_c = 1 \mid \mathbf{x}_c), \tag{6}$$

where $y_c \in \{0, 1\}$ indicates (with $y_c = 1$) that client c should be reported for money laundering. Note that suspicious behavior flagging and client risk profiling can overlap at the client level. Indeed, we could use $F(\mathbf{x}_c)$ as a risk profile for client c .

III. CLIENT RISK PROFILING

We find that studies on client risk profiling are characterized by diagnostics, i.e., efforts to find and explain risk factors. Specifically, unsupervised methods are used to search for new “risky” observations or risk factors. On the other hand, supervised methods are used with an explanatory focus. We also find that studies employing unsupervised methods generally use relatively large data sets. By contrast, studies employing supervised methods use small (labeled) data sets. This difference is likely associated with the cost of labeling observations. Finally, we note that while all studies use private data sets, most share a fair amount of information about the features that they use. As we shall see later, this contrasts with the literature on suspicious behavior flagging.

A. UNSUPERVISED CLIENT RISK PROFILING

Alexandre and Balsa [24] employ K -means clustering [25] to construct risk profiles. The algorithm seeks a clustering $\rho : \mathbb{R}^d \rightarrow \{S_1, \dots, S_K\}$ that assigns every client c to a cluster $k = 1, \dots, K$. This is achieved by solving for

$$\{\mu_k\} = \arg \min_{\{\mu_k\}} \sum_{k=1}^K \sum_{\rho_k} \|\mathbf{x}_c - \mu_k\|^2, \tag{7}$$

where $\mu_k \in \mathbb{R}^d$ denotes the mean of cluster k and $\rho_k = \{c = 1, \dots, n \mid \rho(\mathbf{x}_c) = k\}$ denotes the set of clients assigned to cluster k . The problem is addressed in a greedy optimization fashion; iteratively setting $\mu_k = \frac{1}{|\rho_k|} \sum_{\rho_k} \mathbf{x}_c$ and $\rho(\mathbf{x}_c) = \arg \min_{k=1, \dots, K} \|\mathbf{x}_c - \mu_k\|^2$. To evaluate the approach, the authors employ a data set with approximately

2.4 million clients from an undisclosed financial institution. Disclosed features include the average size and number of transactions. The authors implement $K = 7$ clusters, designating two of them as risky. The first contains clients with many transactions but low transaction values. The second contains clients with older accounts but larger transaction values. Finally, the authors employ decision trees (see Section III-B) to find classification rules that emulate the clusters. The motivation is, presumably, that bank officers find it easier to work with rules than with K -means.

Cao and Do [26] present a similar study, applying clustering with slope [27]. Starting with 8,020 transactions from a Vietnamese bank, the authors first change the level of analysis to individual clients. Features include the sum of in- and outgoing transactions, the number of sending and receiving third parties, and the difference between funds sent and received. The authors then discretize features and build clusters based on cluster histograms’ height-to-width ratios. They finally simulate 25 accounts with money laundering behavior, some easily identifiable in the produced clusters. Much may, however, depend on the nature of the simulations.

Paula et al. [28] use an autoencoder neural network to find outlier Brazilian export firms. Neural networks are directed, acyclic graphs connecting computational units (i.e., neurons) in layers. The output of a feedforward neural network with $l = 1, \dots, L$ layers is given by

$$\begin{aligned} nm(\mathbf{x}_c) &= \phi^{(L)} \left(\left(\dots \left(\phi^{(1)}(\mathbf{x}_c \mathbf{W}^{(1)} + \mathbf{b}^{(1)}) \right) \dots \right) \mathbf{W}^{(L)} + \mathbf{b}^{(L)} \right), \end{aligned} \tag{8}$$

where $\mathbf{W}^{(1)} \in \mathbb{R}^{d \times h_1}, \dots, \mathbf{W}^{(L)} \in \mathbb{R}^{h_{(L-1)} \times h_L}$ are weight matrices, $\mathbf{b}^{(1)} \in \mathbb{R}^{h_1}, \dots, \mathbf{b}^{(L)} \in \mathbb{R}^{h_L}$ are biases, and $\phi^{(1)}, \dots, \phi^{(L)}$ are (non-linear) activation functions. Neural networks are commonly trained with iterative gradient-based optimization. This includes backpropagation [29] coupled with stochastic gradient descent [30] or more recent adaptive schemes like Adam [31]. The aim is to minimize a loss function $l(o_c, nm(\mathbf{x}_c))$ over all observations $c = 1, \dots, n$ where o_c is a target value or vector. Autoencoders, as employed by the authors, are a special type of neural networks that seek a latent representation of their inputs. To this end, they employ an encoder-decoder (i.e., “hourglass”) architecture and try to replicate their inputs in their outputs, i.e., have $o_c = \mathbf{x}_c$. The authors specifically use 5 layers with 18, 6, 3, 6, and 18 neurons. The first two layers (with 18 and 6 neurons) form an encoder. The middle layer with 3 neurons then obtains a latent representation. Finally, the last two layers (with 6 and 18 neurons) form a decoder. The approach is tested on a data set with 819,990 firms. Features include information about debit and credit transactions, export volumes, taxes paid, and previous customs inspections. As a measure of risk, the authors employ the reconstruction error $\rho(\mathbf{x}_c) = \frac{1}{q} \|nm(\mathbf{x}_c) - \mathbf{x}_c\|^2$, frequently used for anomaly or novelty

detection in this setting (see, for instance, [32]). This way, they identify 20 high-risk firms.

B. SUPERVISED CLIENT RISK PROFILING

Colladon and Remondi [33] combine social network analysis and logistic regression. Using 33,670 transactions from an Italian factoring firm, the authors first construct three graphs; \mathcal{G}_1 , \mathcal{G}_2 , and \mathcal{G}_3 . All share the same nodes, representing clients, while edges represent transactions. In \mathcal{G}_1 , edges are weighted relative to transaction size. In \mathcal{G}_2 , they are weighted relative to connected clients' business sectors. Finally, in \mathcal{G}_3 , they are weighted relative to geographic factors. Next, a set of graph metrics are used to construct features for every client. These include in-, out-, and total-degrees, closeness, betweenness, and constraint. A label $y_c \in \{0, 1\}$ is also collected for 288 clients, denoting (with $y_c = 1$) if the client can be connected to a money laundering trial. The authors then employ a logistic regression model

$$P(y_c = 1 | \mathbf{x}_c) = \frac{\exp(\boldsymbol{\beta}^\top \mathbf{x}_c)}{1 + \exp(\boldsymbol{\beta}^\top \mathbf{x}_c)}, \quad (9)$$

where $\boldsymbol{\beta} \in \mathbb{R}^d$ denotes the learnable coefficients. The approach achieves an impressive performance. Results indicate that in-degrees over \mathcal{G}_3 and total-degrees over \mathcal{G}_1 are associated with higher risk. By contrast, constraint over \mathcal{G}_2 and closeness over \mathcal{G}_1 are associated with lower risk.

Rambharat and Tschirhart [34] use panel data from a financial institution in the United States. The data tracks risk profiles $y_{c_p} \in \{1, 2, 3, 4\}$, assigned to $c = 1, \dots, 494$ clients over $p = 1, \dots, 13$ periods. Specifically, y_{c_p} represents low-, medium-, and two types of high-risk profiles. Period-specific features $\mathbf{x}_{c_p} \in \mathbb{R}^d$ include information about clients' business departments, four non-specified "law enforcement actions", and dummy (one-hot encoded) variables that capture the time dimension. To model the data, the authors use an ordinal random effects model where errors and fixed effects are assumed to follow Gaussian distributions. If we let $\Phi(\cdot)$ denote the standard Gaussian cumulative distribution function, the model can be expressed as

$$P(y_{c_p} = m | \mathbf{x}_{c_p}, \alpha_c, \boldsymbol{\beta}, \theta_m, \theta_{m-1}, \sigma_\alpha) = \Phi(\theta_m - \boldsymbol{\beta}^\top \mathbf{x}_{c_p} - \alpha_c) - \Phi(\theta_{m-1} - \boldsymbol{\beta}^\top \mathbf{x}_{c_p} - \alpha_c), \quad (10)$$

where α_c denotes a random client effect, $\boldsymbol{\beta} \in \mathbb{R}^d$ denotes coefficients, and θ_m represents a cut-off value transforming a continuous latent variable $y_{c_p}^*$ into y_{c_p} . Specifically, we have $y_{c_p} = m$ if and only if $\theta_{m-1} < y_{c_p}^* \leq \theta_m$. The level of confidentiality makes it hard to generalize results from the study. The study does, however, illustrate that banks can benefit from a granular risk rating of high-risk clients.

Martínez-Sánchez et al. [35] use decision trees to model clients of a Mexican financial institution. Decision trees [36] are flowchart-like models where internal nodes split the feature space into mutually exclusive sub-regions. Final nodes, called leaves, label observations using a voting system. The authors use data on 181 clients, all labeled as either high-risk or low-risk. Features include information about seniority,

residence, and economic activity. Notably, no train-test split is used. This makes the focus on diagnostics apparent. The authors find that clients with more seniority are comparatively riskier.

Badal-Valero et al. [37] combine Benford's Law and four machine learning models. Benford's Law [38] gives an empirical distribution of leading digits. The authors use it to extract features from financial statements. Specifically, they consider statements from 335 suppliers to a company on trial for money laundering. Of these, 23 suppliers have been investigated and labeled as colluders. All other (non-investigated) suppliers are treated as benevolent. The motivating idea is that any colluders, hiding in the non-investigated group, should be misclassified by the employed models. These include a logistic regression, feedforward neural network, decision tree, and random forest. Random forests [39], in particular, combine multiple decision trees. Every tree uses a random subset of features in every node split. To address class imbalance, i.e., the unequal distribution of labels, the authors investigate weighting and synthetic minority oversampling [40]. The former weighs observations during training, giving higher importance to data from the minority class. The latter balances the data before training, generating synthetic observations of the minority class. According to the authors, synthetic minority oversampling works the best. However, the conclusion is apparently based on simulated evaluation data.

González and Valásquez [41] employ a decision tree, feedforward neural network, and Bayesian network to model Chilean firms using false invoices. Bayesian networks [42], in particular, are probabilistic models that represent variable dependencies via directed acyclic graphs. The authors use data on 582,161 firms, 1,692 of which have been labeled as either fraudulent or non-fraudulent. Features include information about previous audits and taxes paid. Because most firms are unlabeled, the authors first use unsupervised learning to characterize high-risk behavior. To this end, they employ self-organizing maps [43] and neural gas [44]. Both are neural network techniques that build on competitive learning [45] rather than error correction (i.e., gradient-based optimization). While the methods do produce clusters with some behavioral patterns, they do not appear useful for false invoice detection. On the labeled training data, the feedforward neural network achieves the best performance.

IV. SUSPICIOUS BEHAVIOR FLAGGING

We find that the literature on suspicious behavior flagging is characterized by a large proportion of short and suggestive papers. This includes applications of fuzzy logic [46], autoregression [47], and sequence matching [48]. Very few studies apply outlier or anomaly detection techniques [12]. In contrast to work by Canhoto [49], our review demonstrates that there is ample scope to employ both unsupervised and supervised methods for suspicious behavior flagging. Studies using unsupervised methods, however, often contain little performance evaluation. By contrast, studies that use

supervised methods naturally use (a part of) their labeled data for evaluation. In line with thoughts by Breiman [50] (on fraud detection), there is some evidence that supervised methods might perform better than unsupervised methods; see the last part of Section IV-B. However, different types of employed data and the small size of the literature make it difficult to draw a conclusion. Furthermore, non-disclosed features and hand-crafted risk indices generally make it difficult to compare studies.

A. UNSUPERVISED SUSPICIOUS BEHAVIOR FLAGGING

Larik and Haider [51] flag transactions with a combination of principal component analysis and K -means. Given data on approximately 8.2 million transactions, the authors first seek to cluster clients. To this end, principal component analysis [52] is applied to client features $\mathbf{x}_c \in \mathbb{R}^d$, $c = 1, \dots, n$. The method seeks lower-dimensional, linear transformations $\mathbf{z}_c \in \mathbb{R}^q$, $q < d$, that preserve the greatest amount of variance. Let \mathbf{S} denote the data covariance matrix. The first coordinate of \mathbf{z}_c , called the first principal component, is then given by $\mathbf{u}_1^T \mathbf{x}_c$ where the principal direction $\mathbf{u}_1 \in \mathbb{R}^d$ is determined by

$$\begin{aligned} \mathbf{u}_1 = \arg \max_{\mathbf{u} \in \mathbb{R}^d} \quad & \mathbf{u}^T \mathbf{S} \mathbf{u} \\ \text{s.t.} \quad & \mathbf{u}^T \mathbf{u} = 1 \end{aligned} \quad (11)$$

By analogy, the j 'th principal component is given by $\mathbf{u}_j^T \mathbf{x}_c$ where $\mathbf{u}_j \in \mathbb{R}^d$ maximizes $\mathbf{u}_j^T \mathbf{S} \mathbf{u}_j$ subject to $\mathbf{u}_j^T \mathbf{u}_j = 1$ and orthogonality with the previous principal components $h = 1, \dots, j-1$. Principal components are commonly obtained by the eigenvectors of \mathbf{S} corresponding to maximal eigenvalues. Next, the authors use a modified version of K -means to cluster \mathbf{z}_c , $c = 1, \dots, n$. The modification introduces a parameter to control the maximum distance between an observation and the mean of its assigned cluster. A hand-crafted risk index is then used to score and flag incoming transactions. The index compares the sizes and frequencies of transactions within assigned client clusters. As no labels are available, evaluation is limited.

Rocha-Salazar et al. [53] mix fuzzy logic, clustering, and principal component analysis to raise alarms. With fuzzy logic [54], experts first assign risk scores to feature values. These include information about client age, nationality, and transaction statistics. Next, strict competitive learning, fuzzy C -means, self-organizing maps, and neural gas are used to build client clusters. The authors find that fuzzy C -means [55], in particular, produces the best clusters. This algorithm is similar to K -means but uses scores to express degrees of cluster membership rather than hard assignments. The authors further identify one high-risk cluster. Transactions in this cluster are then scored with a hand-crafted risk index. This builds on principal component analysis, weighing features relative to their variances. Data from a Mexican financial institution is used to evaluate the approach. Training is done with 26,751 private and 3,572 business transactions; testing with 1,000 private and 600 business transactions. The

approach shows good results on balanced accuracy (i.e., the average of the true positive and true negative rate).

Raza and Haider [56] propose a combination of clustering and dynamic Bayesian networks. First, client features \mathbf{x}_c are clustered with fuzzy C -means. For each cluster, a q -step dynamic Bayesian network [57] is then trained on transaction sequences $\mathcal{X}_{(c,a)} = \{\mathbf{x}_{(c,a,1)}, \dots, \mathbf{x}_{(c,a,T(a,c))}\}$. Transaction features $\mathbf{x}_{c,a,t}$ include information about amount, period, and type. At test time, incoming transactions (along with the previous $q = 1, 2$ transactions) are passed through the network. A hand-crafted risk index, building on outputted posterior probabilities, is then calculated. The approach is implemented on a data set with approximately 8.2 million transactions (presumably the same data used by Larik and Haider [51]). However, as no labels are available, evaluation is limited.

Camino et al. [58] flag clients with three outlier detection techniques: an isolation forest, a one-class support vector machine, and a Gaussian mixture model. Isolation forests [59] build multiple decision trees using random feature splits. Observations isolated by comparatively few feature splits (averaged over all trees) are then considered outliers. One-class support vector machines [60] use a kernel function to map data into a reproducing Hilbert space. The method then seeks a maximum margin hyperplane that separates data points from the origin. A small number of observations are allowed to violate the hyperplane; these are considered outliers. Finally, Gaussian mixture models [61] assume that all observations are generated by a number of Gaussian distributions. Observations in low-density regions are then considered outliers. The authors combine all three techniques into a single ensemble method. The method is tested on a data set from an AML software company. This contains one million transactions with client-level features recording summary statistics. The authors report positive feedback from the data-supplying company; otherwise, evaluation is limited.

Sun et al. [62] apply extreme value theory [63] to flag outliers in transaction streams. The authors start by engineering two features. The first records the number of times an account has reached a balanced state, i.e., when money transferred into an account is transferred out again. The second records the number of effective fan-ins associated with an account, i.e., when money transferred into the account surpasses a given limit and the account again reaches a balanced state. Next, the Pickands–Balkema–De Haan theorem [64], [65] is invoked to model (derived) conditional feature exceedances according to a generalized Pareto distribution. The approach allows the authors to flag transactions according to a probabilistic limit p (in analogy to the p -values used to test null hypotheses). The approach is tested on real bank data with simulated noise and outliers.

B. SUPERVISED SUSPICIOUS BEHAVIOR FLAGGING

Deng et al. [66] combine logistic regression, stochastic approximation, and sequential D -optimal design for active

learning. The question is how we sequentially should select new observations for inquiry (revealing $y_{(c,a)}$) and use them in the estimation of

$$F(\mathbf{x}_{(c,a)}) = P(y_{(c,a)} = 1 \mid \mathbf{x}_{(c,a)}). \quad (12)$$

The authors employ a data set with 92 inquired accounts and two highly engineered features. The first feature $x_{(c,a)}^{(1)} \in \mathbb{R}$ captures the velocity and size of transactions; the second $x_{(c,a)}^{(2)} \in \mathbb{R}$ captures peer comparisons. Assuming that $F(\cdot)$ is an increasing function in both features, the authors further define a synthetic variable $z_{(c,a)} = \omega x_{(c,a)}^{(1)} + (1 - \omega)x_{(c,a)}^{(2)}$ for $\omega \in [0, 1]$. Finally, $z_{(c,a)}$ is subject to a univariate logistic regression on $y_{(c,a)}$. This allows a combination of stochastic approximation [67] and sequential D -optimal design [68] for new observation selection. The approach significantly outperforms random selection. Furthermore, simulations show that it is robust to underlying data distributions.

Borrajó et al. [69] argue that AML models may benefit from other types of information than simple transaction statistics. To this end, the authors consider behavior trances. These, among other things, contain information about account creation and company ownership. Using custom distance functions, the authors apply K -nearest neighbors [70], [71] to flag illicit behavior. The method predicts that a new observation belongs to the same class as the majority of its k nearest neighbors. While the authors report excellent results, these are, notably, obtained on simulated data.

Zhang and Trubey [72] employ six machine learning models to predict the outcome of AML alarm inquiries. We note that the setup can be used both to qualify existing and raise new alarms under appropriate assumptions. Indeed, let $s_c \in \{0, 1\}$ indicate (with $s_c = 1$) that client c is flagged by a traditional AML system. Assuming that s_c and y_c are conditionally independent given \mathbf{x}_c , we have that

$$P(y_c = 1 \mid \mathbf{x}_c) = P(y_c = 1 \mid \mathbf{x}_c, s_c = 1). \quad (13)$$

If we also assume that $P(s_c = 1 \mid \mathbf{x}_c) > 0$ for all $\mathbf{x}_c \in \{\mathbf{x} \in \mathbb{R}^d : P(y_c = 1 \mid \mathbf{x}) > 0\}$, we can use a model, only trained on previously flagged clients, to raise new alarms. The authors use a data set with 6,113 alarms from a financial institution in the United States. Of these, 34 alarms were reported to authorities. The data set contains ten non-disclosed features. In order to address class imbalance, the authors investigate random over- and undersampling. Both techniques, in particular, increase the performance of a support vector machine [73]. This model seeks to maximize the margin between feature observations of the two classes and a class separating hyperplane (possibly in transformed space). However, a feedforward neural network, robust to both sampling techniques, shows the best performance.

Jullum et al. [74] use gradient boosted trees to model AML alarms. The approach additively combines f_1, \dots, f_K regression trees (i.e., decision trees with continuous outputs) and is implemented with XGBoost [75]. Data comes from a Norwegian bank and contains:

- 1) 16,192 non-flagged transactions,
- 2) 14,932 flagged transactions, dismissed after brief inquiries,
- 3) 1260 flagged transactions, thoroughly inquired but dismissed, and
- 4) 750 flagged and reported transactions.

The authors primarily perform binary classification. Here, transactions in (1)-(3) are treated as licit; transactions in (4) are treated as illicit. Features include information about client background, behavior, and previous AML alarms. To compare model performance with traditional AML systems, the authors propose an evaluation metric called ‘‘proportion of positive predictions’’ (PPP). This records the proportion of positive predictions when classification thresholds are adjusted to obtain a pre-specified true positive rate. Results, in particular, indicate that the inclusion of type (1) transactions improve performance.

Tertychnyi et al. [76] propose a two-layer approach to flag suspicious clients. In the first layer, a logistic regression is used to filter out clients with transaction patterns that are clearly non-illicit. In the second layer, the remaining clients are subject to gradient boosted trees implemented with CatBoost [77]. The authors employ a data set from an undisclosed bank. This contains approximately 330,000 clients from three countries. About 0.004% of the clients have been reported for money laundering. The remaining are randomly sampled. Client-level features include demographic data and transaction statistics. Model performance varies significantly over the three countries in the authors’ data set. However, the performance decreases when each country is modeled separately.

Eddin et al. [78] investigate how aggregated transaction statistics and different graph features can be used to flag suspicious bank client behavior. To this end, the authors consider a random forest model, generalized linear model [79], and gradient boosted trees with LightGBM [80]. The authors utilize a large data set from a non-disclosed bank. This contains 500,000 flagged transactions distributed over 400,000 accounts (3% of which are deemed truly suspicious and labeled as positives). To construct graph features on the data, the authors treat accounts as nodes and transactions as directed edges. Results indicate that the inclusion of Guilty-Walker features [81], using random walks to capture the distances between a given node and illicit nodes, increases model performance.

Charitou et al. [82] combine a sparse autoencoder and a generative adversarial network to flag money laundering in online gambling. The sparse autoencoder is first used to obtain higher-dimensional latent feature encodings. The goal is to increase the distance between positive (i.e., illicit) and negative (i.e., licit) observations. The latent encodings are then used to train a generative adversarial network [83]. This is composed of two competing networks. A generative network produces synthetic observations from Gaussian noise. A discriminative network tries to separate these from real observations and determine the class of the observations. The

approach is tested on multiple data sets. In an AML context, the most relevant of these pertains to money laundering in online gambling. This data set contains 4,700 observations (1,200 of which were flagged for potential money laundering). fi Weber et al. [84] use graph convolutional neural networks to flag suspicious bitcoin transactions. An open data set is provided by Elliptic, a private cryptocurrency analytics company. The data set contains a transaction graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $|\mathcal{V}| = 203,769$ nodes and $|\mathcal{E}| = 234,355$ edges. Nodes represent bitcoin transactions, while edges represent directed payment flows. Using a heuristic approach, 21% of the nodes are labeled as licit; 2% as illicit. For all nodes, 166 features are recorded. Of these, 94 record local information, while the remaining 72 record one-hop information. Graph convolutional neural networks [85] are neural networks designed to work on graph data. Let $\hat{\mathbf{A}}$ denote the normalized adjacency matrix of graph \mathcal{G} . The output of the network's l 'th layer is obtained by

$$\mathbf{H}^{(l)} = \phi^{(l)} \left(\hat{\mathbf{A}} \mathbf{H}^{(l-1)} \mathbf{W}^{(l)} \right), \quad (14)$$

where $\mathbf{W}^{(l)} \in \mathbb{R}^{h_{(l-1)} \times h_{(l)}}$ is a weight matrix, $\mathbf{H}^{(l-1)} \in \mathbb{R}^{|\mathcal{V}| \times h_{(l-1)}}$ is the output from layer $l-1$ (initiated with feature values), and $\phi^{(l)}$ is an activation function. While the best performance is achieved by a random forest model, the graph convolutional neural network proved competitive. Utilizing a time dimension in the data, the authors also fit a temporal graph convolutional neural network [86]. This outperforms the simple graph convolutional neural network. However, it still falls short of the random forest model.

We finally highlight three recent studies that use the Elliptic data set [84]. Alarab et al. [87] propose a neural network structure where graph convolutional embeddings are concatenated with linear embeddings of the original features. This increases model performance significantly. Vassallo et al. [88] investigate the use of gradient boosting on the Elliptic data. Results, in particular, indicate that gradient boosted trees outperform random forests. Furthermore, the authors propose an adapted version of XGBoost to reduce the impact of concept drift. Lorenz et al. [89] experiment with unsupervised anomaly detection. The authors try seven different techniques: local outlier factor [90], K -nearest neighbors [70], [71], principal component analysis [52], one-class support vector machine [60], cluster-based outlier factor [91], angle-based outlier detection [92], and isolation forest [59]. For evaluation, the F1-score is used, recording the harmonic mean between precision and recall (i.e., the true positive rate). Strikingly, all seven unsupervised methods perform substantially worse than a supervised random forest benchmark. As noted by the authors, this contradicts previous literature on unsupervised behavior flagging (see, for example, [58]). One possible explanation is that the Elliptic data, constructed over bitcoin transactions, is qualitatively different from bank transaction data. The authors, following Deng et al. [66], further experiment with four active learning strategies combined with a random forest, gradient boosted trees, and logistic

regression model. Two of the active learning strategies build on unsupervised techniques: elliptic envelope [93] and isolation forest [59]. The remaining two build on supervised techniques: uncertainty sampling [94] and expected model change [95]. Results show that the supervised techniques perform the best.

V. FUTURE RESEARCH DIRECTIONS

Our review reveals that class imbalance and the lack of publicly available data sets are central challenges to AML research. Both may motivate the use of synthetic data. We also note how banks hold vast amounts of high-dimensional and unlabeled data [96]. This may motivate the use of dimension reduction and semi-supervised learning techniques. Other possible research directions include data visualization, deep learning, and interpretable and fair machine learning. In the following, we introduce each of these topics. We also provide brief descriptions of related methods and techniques within each topic.

A. CLASS IMBALANCE, EVALUATION METRICS, AND SYNTHETIC DATA

Due to class imbalance, AML systems tend to label all observations as benevolent. This implies that accuracy is a poor evaluation metric. Instead, we highlight the receiver operating characteristic (ROC) curve [97], plotting true positive versus false positive rates for varying classification thresholds. The area under a ROC curve, called ROCAUC (or sometimes just AUC), is a measure of separability; equal to 1 for perfect classifiers and 0.5 for naive classifiers. Another possible evaluation tool is the precision-recall (PR) [98] curve, plotting precision versus true positive rates for varying classification thresholds. This curve is particularly relevant when class imbalance is severe and true positive rates are of high importance. Notably, both ROC and PR curves consider the relative ranking of predictions for binary outcome models. For multi-class models, Cohen's κ [99] is appealing. This metric evaluates the agreement between two labelings, accounting for agreement by chance. Finally, note that none of the presented metrics introduced above consider calibration, i.e., if model outputs reflect true likelihoods.

To combat class imbalance, data augmentation can be used. Simple approaches include under- and oversampling (see, for instance, [100]). Synthetic minority oversampling (SMOTE) by Chawla et al. [40] is another option for vector data. The technique generates convex combinations of minority class observations. Extensions include borderline-SMOTE [101] and borderline-SMOTE-SVM [102]. These generate observations along estimated decision boundaries. Another SMOTE variant, ADASYN [103], generates observations according to data densities. For time series data (e.g., transaction sequences), there is relatively little literature on data augmentation [104]. Some basic transformations are:

- 1) window cropping, where random time series slices are extracted,

- 2) window wrapping, compressing (i.e., down-sampling), or extending (i.e., up-sampling) time series slices,
- 3) flipping, where the signs of time series are flipped (i.e., multiplied with -1), and
- 4) noise injection, where (typically Gaussian) noise is added to time series.

A few advanced methods also bear mentioning. Teng et al. [105] propose a wavelet transformation to preserve low-frequency time series patterns while noise is added to high-frequency patterns. Iwana and Uchida [106] utilize the element alignment properties of dynamic time wrapping to mix patterns; features of sample patterns are wrapped to match the time steps of reference patterns. Finally, some approaches combine multiple transformations. Cubuk et al. [107] propose to combine transformations at random. Fons et al. [108] propose two adaptive schemes; the first weighs transformed observations relative to a model's loss, the second selects a subset of transformations based on rankings of prediction losses.

Simulating known or hypothesized money laundering patterns from scratch may be the only option for researchers with no available data. Used together with private data sets, the approach may also ensure some reproducibility and generalizability. We refer to the work by Lopez-Rojas and Axelson [109] for an in-depth discussion of simulated data for AML research. The authors develop a simulator, PaySim, for mobile phone transfers. The simulator is, in particular, employed by [110], proposing a generalized version of Isolation Forests to flag suspicious transactions. Weber et al. [111] and Suzumura and Kanezashi [112] further augment PaySim, tailoring it to a more classic bank setting.

We have found only one public data set within the AML literature: the Elliptic data set [84]. This contains a graph over bitcoin transactions. We do, however, note that graph-based approaches may be difficult to implement in a bank setting. Indeed, any bank only knows about transactions going to or from its own clients. Instead, graph approaches may be more relevant for authorities' treatment of AML reports; see work by Savage et al. [19], Drezewski et al. [20], Li et al. [21], and Baltoi et al. [22].

B. VISUALIZATION, DIMENSION REDUCTION, AND SEMI-SUPERVISED LEARNING

Visualization techniques may help identify money laundering [113]. One option is t -distributed stochastic neighbor embedding [114] and its parametric counterpart [115]. The approach is often used for 2- or 3-dimensional embeddings, aiming to keep similar observations close and dissimilar observations distant. First, a probability distribution over pairs of observations is created in the original feature space. Here, similar observations are given higher probability; dissimilar observations are given lower. Next, we seek projections that minimize the Kullback-Leibler divergence [116] to a distribution in a lower-dimensional space. Another option is ISOMAP [117]. This extends multidimensional

scaling [118], using the shortest path between observations to capture intrinsic similarity.

Autoencoders, as discussed in Section III-A, can be used both for dimension reduction, synthetic data generation, and semi-supervised learning. The latter is relevant when we have data sets with many unlabeled (but also some labeled) observations. Indeed, we may train an autoencoder with all the observations. Lower layers can then be reused in a network trained to classify labeled observations. A seminal type of autoencoder was proposed by Kingma and Welling [119]: the variational autoencoder. This is a probabilistic, generative model that seeks to minimize a loss function with two parts. The first part employs the normal reconstruction error. The second part employs the Kullback-Leibler divergence to push latent feature representations toward a Gaussian distribution. An extension, conditional variational autoencoders [120] take class labels into account, modeling a conditional latent variable distribution. This allows us to generate class specific observations. Generative adversarial networks [83] are another option. Here, two neural networks compete against each other; a generative network produces synthetic observations, while a discriminative network tries to separate these from real observations. In analogy with conditional variational autoencoders, conditional generative adversarial nets [121] take class labels into account. Specifically, class labels are fed as inputs to both the discriminator and generator. This may, again, be used to generate class specific observations. While most generative adversarial network methods have been designed to work with visual data, methods applicable to time-series data have recently been proposed [122], [123].

C. NEURAL NETWORKS, DEEP LEARNING, AND TRANSFER LEARNING

The neural networks used in current AML research are generally small and shallow. Deep neural networks, by contrast, employ multiple layers. The motivating idea is to derive higher-level features directly from data. This has, in particular, proved successful for computer vision [124], [125], natural language processing [126], [127], and high-frequency financial time-series analysis [128], [129], [130]. Some authors have also proposed to use the approach to check KYC image information (e.g., driver's licenses) [131], [132] or ease alarm inquiries with sentiment analysis [133].

State-of-the-art deep neural networks use multiple methods to combat unstable gradients. This includes rectified [134], [135] and exponential [136], [137] linear units. Weight initialization is done with Glorot and Bengio [138], He et al. [139], or LeCun et al. [140] initialization. Batch normalization [141] is used to standardize, re-scale, and shift inputs. For recurrent neural networks (introduced below), gradient clipping [142] and layer normalization [143] are often used. Finally, residual or skip connections [144], [145] feed intermediate outputs multiple levels up a network hierarchy.

State-of-the-art networks also use regularization techniques to combat overfitting. Dropout [146], [147] temporarily removes neurons during training, forcing non-dropped neurons to capture more robust relationships. Regularization [148] limits network weights by adding penalty terms to a model's loss function. Finally, max-norm regularization [149] restricts network weights directly during training.

Multiple deep learning methods have been proposed for transfer learning. We refer to the work by Weiss et al. [150] for an extensive review. The general idea is to utilize knowledge across different domains or tasks. One common approach starts by training a neural network on some source problem. Weights (usually from lower layers) are subsequently transferred to a new neural network that is fine-tuned (i.e., re-trained) on another target problem. This may work well when the first neural network learns to extract features that are relevant to both the source and target problem [151]. A subcategory of transfer learning, domain adaptation explicitly tries to alleviate distributional differences across domains. To this end, both unsupervised and supervised methods may be employed (depending on whether or not labeled target data is available). For example, Ganin and Lempitsky [152] propose an unsupervised technique that employs a gradient reversal layer and backpropagation to learn shift invariant features. Tzeng et al. [153] consider a semi-supervised setup where little labeled target data is available. With unlabeled target data, the authors first optimize feature representations to minimize the distance between a source and target distribution. Next, a few labeled target observations are used as reference points to adjust similarity structures among label categories. Finally, we refer to the work by Hedegaard et al. [154] for a discussion and critique of the generic test setup used in the supervised domain adaptation literature and a proposal of a fair evaluation protocol.

Deep neural networks can, like their shallow counterparts, model sequential data. Here, we provide brief descriptions of simple instantiations of such networks. We use the notation introduced in Section III-A and only describe single layers. To form deep learning models, one stacks multiple layers; each layer receives as input the output of its predecessor. Parameters (across all layers) are then jointly optimized by an iterative optimization scheme, as described in Section III-A. Recurrent neural networks are one approach to modeling sequential data. Let $\mathbf{x}_{(t)} \in \mathbb{R}^d$ denote some layer input at time $t = 1, \dots, T$. We can describe the time t output of a basic recurrent neural network layer with m neurons by

$$\mathbf{y}_{(t)} = \phi(\mathbf{W}_x^\top \mathbf{x}_{(t)} + \mathbf{W}_y^\top \mathbf{y}_{(t-1)} + \mathbf{b}), \quad (15)$$

where $\mathbf{W}_x \in \mathbb{R}^{d \times m}$ is an input weight matrix, $\mathbf{W}_y \in \mathbb{R}^{m \times m}$ is an output weight matrix, $\mathbf{b} \in \mathbb{R}^m$ is a bias vector, and $\phi(\cdot)$ is an activation function. Advanced architectures use gates to regulate the flow of information. Long short-term memory (LSTM) cells [155], [156], [157] are one option. Let \odot denote the Hadamard product and σ the standard sigmoid function. At time t , an LSTM layer with m neurons is described by

- 1) an input gate $\mathbf{i}_{(t)} = \sigma(\mathbf{W}_{x,i}^\top \mathbf{x}_{(t)} + \mathbf{W}_{y,i}^\top \mathbf{y}_{(t-1)} + \mathbf{b}_i)$,
- 2) a forget gate $\mathbf{f}_{(t)} = \sigma(\mathbf{W}_{x,f}^\top \mathbf{x}_{(t)} + \mathbf{W}_{y,f}^\top \mathbf{y}_{(t-1)} + \mathbf{b}_f)$,
- 3) an output gate $\mathbf{o}_{(t)} = \sigma(\mathbf{W}_{x,o}^\top \mathbf{x}_{(t)} + \mathbf{W}_{y,o}^\top \mathbf{y}_{(t-1)} + \mathbf{b}_o)$,
- 4) a main transformation
 $\mathbf{g}_{(t)} = \tanh(\mathbf{W}_{x,g}^\top \mathbf{x}_{(t)} + \mathbf{W}_{y,g}^\top \mathbf{y}_{(t-1)} + \mathbf{b}_g)$,
- 5) a long-term state $\mathbf{l}_{(t)} = \mathbf{f}_{(t)} \odot \mathbf{l}_{(t-1)} + \mathbf{i}_{(t)} \odot \mathbf{g}_{(t)}$, and
- 6) an output $\mathbf{y}_{(t)} = \mathbf{o}_{(t)} \odot \tanh(\mathbf{l}_{(t)})$,

where $\mathbf{W}_{x,i}$, $\mathbf{W}_{x,f}$, $\mathbf{W}_{x,g}$, and $\mathbf{W}_{x,o}$ in $\mathbb{R}^{d \times m}$ denote input weight matrices, $\mathbf{W}_{y,i}$, $\mathbf{W}_{y,f}$, $\mathbf{W}_{y,g}$, and $\mathbf{W}_{y,o}$ in $\mathbb{R}^{m \times m}$ denote output weight matrices, and \mathbf{b}_i , \mathbf{b}_f , \mathbf{b}_o , and \mathbf{b}_g in \mathbb{R}^m denote biases. Cho et al. [158] propose a simpler architecture based on gated recurrent units (called GRUs). An alternative to recurrent neural networks, the temporal neural bag-of-features architecture has proved successful for financial time series classification [159]. Here, a radial basis function layer with $k = 1, \dots, K$ neurons is used

$$\rho(\mathbf{x}_{(t)})_k = \frac{\exp(-\|\mathbf{x}_{(t)} - \mathbf{v}_k \odot \mathbf{w}_k\|_2)}{\sum_{k=1}^K \exp(-\|\mathbf{x}_{(t)} - \mathbf{v}_k \odot \mathbf{w}_k\|_2)}, \quad (16)$$

where \mathbf{v}_k and \mathbf{w}_k in \mathbb{R}^d are weights that describe the k 'th neuron's center and width, respectively. Next, an accumulation layer is used to find a constant length representation in \mathbb{R}^K ,

$$\mathbf{h} = \frac{1}{T} \sum_{t=1}^T \begin{pmatrix} \rho(\mathbf{x}_{(t)})_1 \\ \vdots \\ \rho(\mathbf{x}_{(t)})_K \end{pmatrix}. \quad (17)$$

Bilinear neural networks may also be used to model time domain information. Let $\mathbf{X} = [\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(T)}]$ be a matrix with columns $\mathbf{x}_{(t)} \in \mathbb{R}^d$ for $t = 1, \dots, T$. A temporal bilinear layer with m neurons can then be described as

$$\mathbf{Y} = \phi(\mathbf{W}_1 \mathbf{X} \mathbf{W}_2 + \mathbf{B}), \quad (18)$$

where $\mathbf{W}_1 \in \mathbb{R}^{m \times d}$ and $\mathbf{W}_2 \in \mathbb{R}^{T \times T'}$ are weight matrices and $\mathbf{B} \in \mathbb{R}^{m \times T'}$ is a bias matrix. Notably, \mathbf{W}_1 models feature interactions at fixed time points while \mathbf{W}_2 models feature changes over time.

Attention mechanisms have recently become state-of-the-art. These allow neural networks to dynamically focus on relevant sequence elements. Bahdanau et al. [160] consider a bidirectional recurrent neural network [161] and propose a mechanism known as additive or concatenative attention. The mechanism assumes an encoder-decoder architecture. During decoding, it computes a context vector by weighing an encoder's hidden states. Weights are obtained by a secondary feedforward neural network (called an alignment model) and normalized by a softmax layer (to obtain attention scores). Notably, the secondary network is trained jointly with the primary network. Luong attention [162] is another popular mechanism, using the dot product between an encoder's and a decoder's hidden states as a similarity measure (the mechanism is also called dot-product attention). Vaswani et al. [163] propose the seminal transformer architecture. Here, an encoder first applies self-attention (i.e., scaled Luong attention). As before, let $\mathbf{X} = [\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(T)}]$

denote our matrix of sequence elements. We can describe a self-attention layer as

$$Z = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}} \right) \mathbf{V}, \quad (19)$$

where

- 1) $\mathbf{Q} \in \mathbb{R}^{T \times d_k}$, called the query matrix, is given by $\mathbf{Q} = \mathbf{X}^\top \mathbf{W}_Q$ with a weight matrix $\mathbf{W}_Q \in \mathbb{R}^{d \times d_k}$,
- 2) $\mathbf{K} \in \mathbb{R}^{T \times d_k}$, called the key matrix, is given by $\mathbf{K} = \mathbf{X}^\top \mathbf{W}_K$ with a weight matrix $\mathbf{W}_K \in \mathbb{R}^{d \times d_k}$, and
- 3) $\mathbf{V} \in \mathbb{R}^{T \times d_v}$, called the value matrix, is given by $\mathbf{V} = \mathbf{X}^\top \mathbf{W}_V$ with a weight matrix $\mathbf{W}_V \in \mathbb{R}^{d \times d_v}$.

Note that the softmax function is applied row-wise. It outputs a $T \times T$ matrix. Here, every row $t = 1, \dots, T$ measures how much attention we pay to $\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(T)}$ in relation to $\mathbf{x}_{(t)}$. During decoding, the transformer also applies self-attention. Here, the key and value matrices are taken from the encoder. In addition, the decoder is only allowed to attend to earlier output sequence elements (future elements are masked, i.e., set to $-\text{inf}$ before softmax is applied). Notably, the authors apply multiple parallel instances of self-attention. The approach, known as multi-head attention, allows attention over many abstract dimensions. Finally, positional encoding, residual connections, layer normalization, and supplementary feedforward layers are used. As a last attention mechanism, we highlight temporal attention augmented bilinear layers [164]. With the notation used to introduce temporal bilinear layers above, we may express a temporal attention augmented bilinear layer as

- 1) $\tilde{\mathbf{X}} = \mathbf{W}_1 \mathbf{X}$,
- 2) $\mathbf{E} = \tilde{\mathbf{X}} \mathbf{W}$,
- 3) $\alpha_{(i,j)} = \frac{\exp(e_{(i,j)})}{\sum_{k=1}^T \exp(e_{(i,k)})}$,
- 4) $\tilde{\mathbf{X}} = \lambda(\tilde{\mathbf{X}} \odot \mathbf{A}) + (1 - \lambda)\tilde{\mathbf{X}}$,
- 5) $\mathbf{Y} = \phi(\tilde{\mathbf{X}} \mathbf{W}_2) + \mathbf{B}$,

where $a_{(i,j)}$ and $e_{(i,j)}$ denote the (i, j) 'th element of $\mathbf{A} \in \mathbb{R}^{m \times T}$ and $\mathbf{E} \in \mathbb{R}^{m \times T}$, respectively, $\mathbf{W} \in \mathbb{R}^{T \times T}$ is a weight matrix with fixed diagonal elements equal to $1/T$, and $\lambda \in [0, 1]$ is a scalar allowing soft attention. In particular, \mathbf{E} is used to express the relative importance of temporal feature instances (learned through \mathbf{W}), while \mathbf{A} contains our attention scores.

D. INTERPRETABLE AND FAIR MACHINE LEARNING

Advanced machine learning models often outperform their simple statistical counterparts. Their behavior can, however, be much harder to understand, interpret, and explain. While some supervisory authorities have shown a fair amount of leeway regarding advanced AML models [165], this is a potential problem.

“Fairness” is an ambiguous concept in machine learning with many different and overlapping definitions [166]. The equalized odds definition states that different protected groups (e.g., genders or races) should have equal true and false positive rates. The conditional statistical parity definition takes a set of legitimate discriminative features into

account, stating that the likelihood of a positive prediction should be the same across protected groups given the set of legitimate discriminative features. Finally, the counterfactual fairness definition is based on a notation that a prediction is fair if it remains unchanged in a counterfactual world where some features of interest are changed. Approaches for fair machine learning also vary greatly. In an exemplary paper, Louizos et al. [167] consider the use of variational autoencoders to ensure fairness. The authors treat sensitive features as nuisance or noise variables, encouraging separation between these and (informative) latent features by using factorized priors and a maximum mean discrepancy penalty term [168]. In another exemplary paper, Zhang et al. [169] propose the use of adversarial learning. Here, a primary model tries to predict an outcome variable while minimizing an adversarial model’s ability to predict protected feature values. Notably, the adversarial model takes as inputs both the primary model’s predictions and other relevant features, depending on the fairness definition of interest.

Regarding interpretability, we follow Du et al. [170] and distinguish between intrinsic and *post-hoc* interpretability. Intrinsically interpretable models are, by design, easy to understand. This includes simple decision trees and linear regression. Notably, attention mechanisms also exhibit some intrinsic interpretability; we may investigate attention scores to see what part of a particular input sequence a neural network focuses on. Other models and architectures work as “black boxes” and require *post-hoc* interpretability methods. Here, it is useful to distinguish between global and local interpretability. The former is concerned with overreaching model behavior; the latter with individual predictions. One possible technique for local interpretability is LIME [171]. Consider a situation where a black box model and a single observation are given. The method then first generates a set of permuted observations (relative to the original observation) with black box model outputs. An intrinsically interpretable model is then trained on the synthetic data. Finally, this model is used to explain the original observation’s black box prediction. Gradient-based methods [172] use the gradients associated with a particular observation and black box model to capture importance. The fundamental idea is that larger gradients (either positive or negative) imply larger feature importance. Individual conditional expectation plots [173] are another option. These illustrate what happens to a particular black box prediction if we vary one feature value of the underlying observation. Similarly, partial dependence plots [174] may be used for global interpretability. Here, we average results from feature variations over all observations in a data set. This may, however, be misleading if input features are highly correlated. In this case, accumulated local effects plots [175] present an attractive alternative. These rely on conditional feature distributions and employ prediction differences. For counterfactual observation generation, numerous methods have been proposed [176], [177], [178]. While these generally need to query an underlying model multiple times, efficient methods utilizing invertible neural networks have

also been proposed [179]. A related problem concerns the quantitative evaluation of counterfactual examples; see the work by Hvilshøj et al. [180] for an in-depth discussion. Finally, we highlight Shapley additive explanations (SHAP) by Lundberg and Lee [181]. The approach is based on Shapley values [182] with a solid game-theoretical foundation. For a given observation, SHAP values record average marginal feature contributions (to a black box model's output) over all possible feature coalitions. The approach allows both local and global interpretability. Indeed, every observation is given a set of SHAP values (one for each input feature). Summed over the entire data set, the (numerical) SHAP values show accumulated feature importance. Although SHAP values are computationally expensive, polynomial time estimation is possible for tree-based models [183].

VI. CONCLUSION

Inspired by FATF's recommendations, we propose a terminology for AML in banks structured around two central tasks: (i) client risk profiling and (ii) suspicious behavior flagging. The former assigns general risk scores to clients (e.g., for use in KYC operations) while the latter raises alarms on clients, accounts, or transactions (e.g., for use in transaction monitoring). Our review reveals that the literature on client risk profiling is characterized by diagnostics, i.e., efforts to find and explain risk factors. The literature on suspicious behavior flagging, on the other hand, is characterized by non-disclosed features and hand-crafted risk indices.

In general, we find that the literature on AML in banks is plagued by a number of problems. Two challenges are class imbalance and a lack of public data sets. To address class imbalance, a multitude of different data augmentation methods may be used. Motivated by the sensitivity of bank data, synthetic data generation may be a viable way to address the lack of public data sets. Synthetic public data sets would, in particular, facilitate better evaluation and reproducibility of, as well as comparisons between, new and existing methods. Other directions for future research include methods for dimension reduction, semi-supervised learning, data visualization, deep learning, and interpretable and fair machine learning. Finally, we strongly advise against the use of accuracy as an evaluation metric for AML applications, instead emphasizing ROC or PR curves.

REFERENCES

- [1] T. Pietschmann, J. Walker, M. Shaw, D. Chryssikos, D. Schantz, P. Davis, C. Philip, A. Korenblik, R. Johansen, S. Kunnen, K. Kuttig, T. L. Pichon, and S. Chawla, "Estimating illicit financial flows resulting from drug trafficking and other transnational organized crimes." United Nations Office Drugs Crime, Vienna, Austria, Tech. Rep., 2011. [Online]. Available: https://www.unodc.org/documents/data-and-analysis/Studies/Illicit_financial_flows_2011_web.pdf
- [2] J. McDowell and G. Novis, "Consequences of money laundering and financial crime," *Econ. Perspect.*, vol. 6, no. 2, pp. 6–8, May 2001.
- [3] J. Ferwerda, "The effects of money laundering," in *Research Handbook on Money Laundering*, B. Unger and D. Linde, Eds. Northampton, MA, USA: Edward Elgar, 2013, pp. 35–46.
- [4] B. L. Bartlett, "The negative effects of money laundering on economic development," *Platypus Mag.*, vol. 77, pp. 18–23, Dec. 2002.
- [5] R. Grint, C. O'Driscoll, and S. Paton, "New technologies and anti-money laundering compliance," Financial Conduct Authority, London, U.K., Tech. Rep., 2017. [Online]. Available: <https://www.fca.org.uk/publication/research/new-technologies-in-aml-final-report.pdf>
- [6] *Opportunities and Challenges of New Technologies for AML/CFT*, Financial Action Task Force, Paris, France, 2021. [Online]. Available: <https://www.fatf-gafi.org/media/fatf/documents/reports/Opportunities-Challenges-of-New-Technologies-for-AML-CFT.pdf>
- [7] The Wolfsberg Group. (2022). *Wolfsberg Principles for Using Artificial Intelligence and Machine Learning in Financial Crime Compliance*. [Online]. Available: <https://www.wolfsberg-principles.com/sites/default/files/wb/Wolfsberg%20Principles%20for%20Using%20Artificial%20Intelligence%20and%20Machine%20Learning%20in%20Financial%20Crime%20Compliance.pdf>
- [8] M. Biallas and F. O'Neill, "Artificial intelligence innovation in financial services," Int. Finance Corp., World Bank Group, Washington, DC, USA, 2020. [Online]. Available: <https://openknowledge.worldbank.org/server/api/core/bitstreams/2edeba59-b334-5ce1-bbdc-fca40a26e08e/content>
- [9] S. Breslow, M. Hagstroem, D. Mikkelsen, and K. Robu. (2017). The new frontier in anti-money laundering. McKinsey and Company. [Online]. Available: https://icsid.worldbank.org/sites/default/files/parties_publications/C6106/2020.01.20%20Claimants_%20Rejoinder/Factual%20Exhibits/C-0901-ENG.pdf
- [10] G. S. Leite, A. B. Albuquerque, and P. R. Pinheiro, "Application of technological solutions in the fight against money laundering—A systematic literature review," *Appl. Sci.*, vol. 9, no. 22, p. 4800, Nov. 2019, doi: 10.3390/app9224800.
- [11] E. W. T. Ngai, Y. Hu, Y. H. Wong, Y. Chen, and X. Sun, "The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature," *Decis. Support Syst.*, vol. 50, no. 3, pp. 559–569, Feb. 2011.
- [12] W. Hilal, S. A. Gadsden, and J. Yawney, "Financial fraud: A review of anomaly detection techniques and recent advances," *Expert Syst. Appl.*, vol. 193, May 2022, Art. no. 116429, doi: 10.1016/j.eswa.2021.116429.
- [13] *International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation*, Financial Action Task Force, Paris, France, Mar. 2022. [Online]. Available: <https://www.fatf-gafi.org/media/fatf/documents/recommendations/pdfs/FATF%20Recommendations%202012.pdf>
- [14] D. Magnusson, "The costs of implementing the anti-money laundering regulations in Sweden," *J. Money Laundering Control*, vol. 12, no. 2, pp. 101–112, May 2009.
- [15] A. Verhage, "Supply and demand: Anti-money laundering by the compliance industry," *J. Money Laundering Control*, vol. 12, no. 4, pp. 371–391, Oct. 2009.
- [16] D. S. Demetis, "Fighting money laundering with technology: A case study of bank X in the U.K.," *Decis. Support Syst.*, vol. 105, pp. 96–107, Jan. 2018.
- [17] B. Richardson, D. Williams, and D. Mikkelsen. (2019). Network analytics and the fight against money laundering. McKinsey and Company. [Online]. Available: <https://www.mckinsey.com/industries/financial-services/our-insights/banking-matters/network-analytics-and-the-fight-against-money-laundering>
- [18] M. Sun, "Suspicious activity reports related to cash transactions surge," *Wall Street J.* Accessed: Dec. 29, 2022. [Online]. Available: <https://www.wsj.com/articles/suspicious-activity-reports-related-to-cash-transactions-surge-11612900800>
- [19] D. Savage, Q. Wang, P. Chou, X. Zhang, and X. Yu, "Detection of money laundering groups using supervised learning in networks," 2016, *arXiv:1608.00708*.
- [20] R. Drezewski, J. Sepielak, and W. Filipkowski, "The application of social network analysis algorithms in a system supporting money laundering detection," *Inf. Sci.*, vol. 295, pp. 18–32, Feb. 2015.
- [21] X. Li, X. Cao, X. Qiu, J. Zhao, and J. Zheng, "Intelligent anti-money laundering solution based upon novel community detection in massive transaction networks on spark," in *Proc. 5th Int. Conf. Adv. Cloud Big Data (CBD)*, Shanghai, China, Aug. 2017, pp. 176–181.
- [22] A. Baltoiu, A. Patrascu, and P. Irofti, "Community-level anomaly detection for anti-money laundering," 2019, *arXiv:1910.11313*.
- [23] R. J. Bolton and D. J. Hand, "Statistical fraud detection: A review," *Stat. Sci.*, vol. 17, no. 3, pp. 235–255, Aug. 2002.

- [24] C. Alexandre and J. Balsa, "Client profiling for an anti-money laundering system," 2015, *arXiv:1510.00878*.
- [25] S. Lloyd, "Least squares quantization in PCM," *IEEE Trans. Inf. Theory*, vol. IT-28, no. 2, pp. 129–137, Mar. 1982.
- [26] D. K. Cao and P. Do, "Applying data mining in money laundering detection for the Vietnamese banking industry," in *Proc. 4th Asian Conf. Intell. Inf. Database Syst.*, Taiwan, Mar. 2012, pp. 207–216.
- [27] Y. Yang, X. Guan, and J. You, "CLOPE: A fast and effective clustering algorithm for transactional data," in *Proc. 8th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Edmonton, AB, Canada, Jul. 2002, pp. 682–687.
- [28] E. L. Paula, M. Ladeira, R. N. Carvalho, and T. Marzagão, "Deep learning anomaly detection as support fraud investigation in Brazilian exports and anti-money laundering," in *Proc. 15th IEEE Int. Conf. Mach. Learn. Appl. (ICMLA)*, Anaheim, CA, USA, Dec. 2016, pp. 954–960.
- [29] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, pp. 533–536, Oct. 1986.
- [30] H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Statist.*, vol. 22, no. 3, pp. 400–407, Sep. 1951.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [32] J. Chen, S. Sathe, C. C. Aggarwal, and D. S. Turaga, "Outlier detection with autoencoder ensembles," in *Proc. SIAM Int. Conf. Data Mining*, Houston, TX, USA, Apr. 2017, pp. 90–98.
- [33] A. F. Colladon and E. Remondi, "Using social network analysis to prevent money laundering," *Expert Syst. Appl.*, vol. 67, pp. 49–58, Jan. 2017.
- [34] B. R. Rambharat and A. J. Tschirhart, "A statistical diagnosis of customer risk ratings in anti-money laundering surveillance," *Stat. Public Policy*, vol. 2, no. 1, pp. 12–24, Apr. 2015.
- [35] J. F. Martínez-Sánchez, S. Cruz-García, and F. Venegas-Martínez, "Money laundering control in Mexico: A risk management approach through regression trees (data mining)," *J. Money Laundering Control*, vol. 23, no. 2, pp. 427–439, Mar. 2020.
- [36] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. Belmont, CA, USA: Wadsworth, 1984.
- [37] E. Badal-Valero, J. A. Alvarez-Jareño, and J. M. Pavia, "Combining Benford's Law and machine learning to detect money laundering. An actual Spanish court case," *Forensic Sci. Int.*, vol. 282, pp. 24–34, Jan. 2018.
- [38] F. Benford, "The law of anomalous numbers," *Proc. Amer. Philos. Soc.*, vol. 78, no. 4, pp. 551–572, Mar. 1938.
- [39] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, Oct. 2001.
- [40] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *J. Artif. Intell. Res.*, vol. 16, pp. 321–357, Jun. 2002.
- [41] P. Castellón González and J. D. Velásquez, "Characterization and detection of taxpayers with false invoices using data mining techniques," *Expert Syst. Appl.*, vol. 40, no. 5, pp. 1427–1436, Apr. 2013.
- [42] J. Pearl, "Bayesian networks: A model of self-activated memory for evidential reasoning," Dept. Comput. Sci., Univ. California, Los Angeles, CA, USA, Tech. Rep. CSD-850021, Jun. 1985. [Online]. Available: https://ftp.cs.ucla.edu/pub/stat_ser/r43-1985.pdf
- [43] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biol. Cybern.*, vol. 43, no. 1, pp. 59–69, Jan. 2004.
- [44] T. Martinetz and K. Schulten, "A 'neural gas' network learns topologies," in *Artificial Neural Network*, T. Kohonen, K. Mäkisara, O. Simula, and J. Kangas, Eds. Amsterdam, The Netherlands: Elsevier, 1991, pp. 397–402. [Online]. Available: <https://www.ks.uiuc.edu/Publications/Papers/PDF/MART91B/MART91B.pdf>
- [45] D. E. Rumelhart and D. Zipser, "Feature discovery by competitive learning," *Cogn. Sci.*, vol. 9, no. 1, pp. 75–112, Jan. 1985.
- [46] Y. Chen and J. Mathe. (2011). *Fuzzy Computing Applications for Anti-Money Laundering and Distributed Storage System Load Monitoring*. [Online]. Available: <https://storage.googleapis.com/pub-tools-public-publication-data/pdf/37118.pdf>
- [47] S. Kannan and K. Somasundaram, "Autoregressive-based outlier algorithm to detect money laundering activities," *J. Money Laundering Control*, vol. 20, no. 2, pp. 190–202, May 2017.
- [48] X. Liu, P. Zhang, and D. Zeng, "Sequence matching for suspicious activity detection in anti-money laundering," in *Proc. IEEE Int. Conf. Intell. Secur. Inform.*, Taiwan, Jun. 2008, pp. 50–61.
- [49] A. I. Canhoto, "Leveraging machine learning in the global fight against money laundering and terrorism financing: An affordances perspective," *J. Bus. Res.*, vol. 131, pp. 441–452, Jul. 2021.
- [50] L. Breiman, "[Statistical fraud detection: A review]: Comment," *Statist. Sci.*, vol. 17, no. 3, pp. 252–254, Aug. 2002.
- [51] A. S. Larik and S. Haider, "Clustering based anomalous transaction reporting," *Proc. Comput. Sci.*, vol. 3, pp. 606–610, Jan. 2011.
- [52] I. T. Jolliffe and J. Cadima, "Principal component analysis: A review and recent developments," *Philos. Trans. Roy. Soc. A, Math., Phys. Eng. Sci.*, vol. 374, no. 2065, Apr. 2016, Art. no. 20150202, doi: [10.1098/rsta.2015.0202](https://doi.org/10.1098/rsta.2015.0202).
- [53] J.-D.-J. Rocha-Salazar, M.-J. Segovia-Vargas, and M.-D.-M. Camacho-Miñano, "Money laundering and terrorism financing detection using neural networks and an abnormality indicator," *Expert Syst. Appl.*, vol. 169, May 2021, Art. no. 114470.
- [54] L. A. Zadeh, "The role of fuzzy logic in the management of uncertainty in expert systems," *Fuzzy Sets Syst.*, vol. 11, nos. 1–3, pp. 199–227, Jul. 1983.
- [55] J. Dunn, "A fuzzy relative of the ISODATA process and its use in detecting compact well-separated clusters," *J. Cybern.*, vol. 3, no. 3, pp. 32–57, 1974.
- [56] S. Raza and S. Haider, "Suspicious activity reporting using dynamic Bayesian networks," *Proc. Comput. Sci.*, vol. 3, pp. 987–991, Jan. 2011.
- [57] P. Dagum, A. Galper, and E. Horvitz, "Dynamic network models for forecasting," in *Proc. 8th Conf. Uncertainty Artif. Intell.*, Stanford, CA, USA, Jul. 1992, pp. 41–48.
- [58] R. D. Camino, R. State, L. Montero, and P. Valtchev, "Finding suspicious activities in financial transactions and distributed ledgers," in *Proc. IEEE Int. Conf. Data Mining Workshops (ICDMW)*, New Orleans, LA, USA, Nov. 2017, pp. 787–796.
- [59] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *Proc. 8th IEEE Int. Conf. Data Mining*, Pisa, Italy, Dec. 2008, pp. 413–422.
- [60] B. Schölkopf, J. Platt, J. C. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural Comput.*, vol. 13, no. 7, pp. 1443–1471, Jul. 2001.
- [61] D. Reynolds, "Gaussian mixture models," in *Encyclopedia of Biometrics*, S. Z. Li and A. K. Jain, Eds. Boston, MA, USA: Springer, 2015, pp. 827–832.
- [62] X. Sun, W. Feng, S. Liu, Y. Xie, S. Bhatia, B. Hooi, W. Wang, and X. Cheng, "MonLAD: Money laundering agents detection in transaction streams," in *Proc. 15th ACM Int. Conf. Web Search Data Mining*, New York, NY, USA, Feb. 2022, pp. 976–986.
- [63] P. Emberchts, C. Klüppelberg, and T. Mikosch, *Modelling Extremal Events*. Berlin, Germany: Springer, 1997.
- [64] J. Pickands, "Statistical inference using extreme order statistics," *Ann. Statist.*, vol. 3, no. 1, pp. 119–131, Jan. 1975.
- [65] A. A. Balkema and L. de Haan, "Residual life time at great age," *Ann. Probab.*, vol. 2, no. 5, pp. 792–804, Oct. 1974.
- [66] X. Deng, V. R. Joseph, A. Sudjianto, and C. F. J. Wu, "Active learning through sequential design, with applications to detection of money laundering," *J. Amer. Stat. Assoc.*, vol. 104, no. 487, pp. 969–981, Sep. 2009.
- [67] C. F. J. Wu, "Efficient sequential designs with binary data," *J. Amer. Stat. Assoc.*, vol. 80, no. 392, pp. 974–984, Dec. 1985.
- [68] B. T. Neyner, "A D -optimality-based sensitivity test," *Technometrics*, vol. 36, no. 1, pp. 61–70, Feb. 1994.
- [69] D. Borrajo, M. Veloso, and S. Shah, "Simulating and classifying behavior in adversarial environments based on action-state traces: An application to money laundering," in *Proc. 1st ACM Int. Conf. AI Finance*, New York, NY, USA, Oct. 2020, Art. no. 3, doi: [10.1145/3383455.3422536](https://doi.org/10.1145/3383455.3422536).
- [70] E. Fix and J. L. Hodges, "Discriminatory analysis. Nonparametric discrimination: Consistency properties," *Int. Stat. Rev.*, vol. 57, no. 3, pp. 238–247, Dec. 1989.
- [71] N. S. Altman, "An introduction to kernel and nearest-neighbor non-parametric regression," *Amer. Statistician*, vol. 46, no. 3, pp. 175–185, Aug. 1992.
- [72] Y. Zhang and P. Trubey, "Machine learning and sampling scheme: An empirical study of money laundering detection," *Comput. Econ.*, vol. 54, no. 3, pp. 1043–1063, Oct. 2019.
- [73] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, Sep. 1995.
- [74] M. Jullum, A. Løland, R. B. Huseby, G. Ånonsen, and J. Lorentzen, "Detecting money laundering transactions with machine learning," *J. Money Laundering Control*, vol. 23, no. 1, pp. 173–186, Jan. 2020.

- [75] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, San Francisco, CA, USA, Aug. 2016, pp. 785–794.
- [76] P. Tertychnyi, I. Slobozhan, M. Ollikainen, and M. Dumas, "Scalable and imbalance-resistant machine learning models for anti-money laundering: A two-layered approach," in *Proc. 10th Int. Workshop Enterprise Appl., Markets Services Finance Ind.*, Helsinki, Finland, Aug. 2020, pp. 43–58.
- [77] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, "CatBoost: Unbiased boosting with categorical features," in *Proc. 32nd Conf. Neural Inf. Process. Syst.*, Montreal, QC, Canada, Dec. 2018, pp. 6639–6649.
- [78] A. N. Eddin, J. Bono, D. Aparício, D. Polido, J. T. Ascensão, P. Bizarro, and P. Ribeiro, "Anti-money laundering alert optimization using machine learning with graphs," 2021, *arXiv:2112.07508*.
- [79] J. A. Nelder and R. W. Wedderburn, "Generalized linear models," *J. Roy. Stat. Soc. A, Gen.*, vol. 135, no. 3, pp. 370–384, 1972.
- [80] G. Ke, Q. Meng, T. Finley, T. Wand, W. Chen, W. Ma, Q. Ye, and T.-Y. Lie, "LightGBM: A highly efficient gradient boosting decision tree," in *Proc. 31st Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Dec. 2017, pp. 3146–3154.
- [81] C. Oliveira, J. Torres, M. I. Silva, D. Aparício, J. T. Ascensão, and P. Bizarro, "GuiltyWalker: Distance to illicit nodes in the Bitcoin network," 2021, *arXiv:2102.05373*.
- [82] C. Charitou, A. D. Garcez, and S. Dragicevic, "Semi-supervised GANs for fraud detection," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2020, pp. 1–8, doi: 10.1109/IJCNN48605.2020.9206844.
- [83] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. 27th Conf. Neural Inf. Process. Syst.*, Montreal, QC, Canada, Dec. 2014, pp. 2672–2680.
- [84] M. Weber, G. Domeniconi, J. Chen, D. K. I. Weidele, C. Bellei, T. Robinson, and C. E. Leiserson, "Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics," 2019, *arXiv:1908.02591*.
- [85] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," 2016, *arXiv:1609.02907*.
- [86] A. Pareja, G. Domeniconi, J. Chen, T. Ma, T. Suzumura, H. Kanezashi, T. Kaler, T. B. Schardl, and C. E. Leiserson, "EvolveGCN: Evolving graph convolutional networks for dynamic graphs," 2019, *arXiv:1902.10191*.
- [87] I. Alarab, S. Prakoonwit, and M. I. Nacer, "Competence of graph convolutional networks for anti-money laundering in Bitcoin blockchain," in *Proc. 5th Int. Conf. Mach. Learn. Technol.*, Beijing, China, Jun. 2020, pp. 23–27.
- [88] D. Vassallo, V. Vella, and J. Ellul, "Application of gradient boosting algorithms for anti-money laundering in cryptocurrencies," *Social Neww. Comput. Sci.*, vol. 2, no. 3, Mar. 2021, Art. no. 143.
- [89] J. Lorenz, M. I. Silva, D. Aparício, J. T. Ascensão, and P. Bizarro, "Machine learning methods to detect money laundering in the Bitcoin blockchain in the presence of label scarcity," in *Proc. 1st ACM Int. Conf. AI Finance*, New York, NY, USA, Oct. 2020, pp. 1–8, doi: 10.1145/3383455.3422549.
- [90] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF: Identifying density-based local outliers," *ACM SIGMOD Rec.*, vol. 29, no. 2, pp. 93–104, Jun. 2000.
- [91] S.-Y. Jiang and Q.-B. An, "Clustering-based outlier detection method," in *Proc. 5th Int. Conf. Fuzzy Syst. Knowl. Discovery*, Jinan, China, Oct. 2008, pp. 429–433.
- [92] H.-P. Kriegel, M. Schubert, and A. Zimek, "Angle-based outlier detection in high-dimensional data," in *Proc. 14th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Las Vegas, NV, USA, Aug. 2008, pp. 444–452.
- [93] P. J. Rousseeuw and K. Van Driessen, "A fast algorithm for the minimum covariance determinant estimator," *Technometrics*, vol. 41, no. 3, pp. 212–223, 1999.
- [94] D. D. Lewis and J. Catlett, "Heterogeneous uncertainty sampling for supervised learning," in *Proc. 11th Int. Conf. Mach. Learn.*, New Brunswick, NJ, USA, Jul. 1994, pp. 148–156.
- [95] B. Settles, M. Craven, and S. Ray, "Multiple-instance active learning," in *Proc. 20th Conf. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2007, pp. 1289–1296.
- [96] A. Sudjianto, M. Yuan, D. Kern, S. Nair, A. Zhang, and F. Cela-Díaz, "Statistical methods for fighting financial crimes," *Technometrics*, vol. 52, no. 1, pp. 5–19, Feb. 2010.
- [97] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, Jun. 2006.
- [98] T. Saito and M. Rehmsmeier, "The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets," *PLoS ONE*, vol. 10, no. 3, Mar. 2015, Art. no. e0118432, doi: 10.1371/journal.pone.0118432.
- [99] M. L. McHugh, "Interrater reliability: The Kappa statistic," *Biochem. Med.*, vol. 22, no. 3, pp. 276–282, Oct. 2012.
- [100] G. Lemaître, F. Nogueira, and C. K. Aridas, "Imbalanced-learn: A Python toolbox to tackle the curse of imbalanced datasets in machine learning," *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 559–563, Jan. 2017.
- [101] H. Han, W.-Y. Wang, and B.-H. Mao, "Borderline-SMOTE: A new over-sampling method in imbalanced data sets learning," in *Proc. Int. Conf. Intell. Comput.*, Hefei, China, Aug. 2005, pp. 878–887.
- [102] H. M. Nguyen, E. W. Cooper, and K. Kamei, "Borderline over-sampling for imbalanced data classification," *Int. J. Knowl. Eng. Soft Data Paradigms*, vol. 3, no. 1, pp. 4–21, Apr. 2011.
- [103] H. He, Y. Bai, E. A. Garcia, and S. Li, "ADASYN: Adaptive synthetic sampling approach for imbalanced learning," in *Proc. IEEE Int. Joint Conf. Neural Netw.*, Hong Kong, Jun. 2008, pp. 1322–1328.
- [104] Q. Wen, L. Sun, F. Yang, X. Song, J. Gao, X. Wang, and H. Xu, "Time series data augmentation for deep learning: A survey," 2020, *arXiv:2002.12478*.
- [105] X. Teng, T. Wang, X. Zhang, L. Lan, and Z. Luo, "Enhancing stock price trend prediction via a time-sensitive data augmentation method," *Complexity*, vol. 2020, pp. 1–8, Feb. 2020, doi: 10.1155/2020/6737951.
- [106] B. K. Iwana and S. Uchida, "Time series data augmentation for neural networks by time warping with a discriminative teacher," 2020, *arXiv:2004.08780*.
- [107] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, "RandAugment: Practical automated data augmentation with a reduced search space," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Seattle, WA, USA, Jun. 2020, pp. 3008–3017.
- [108] E. Fons, P. Dawson, X.-J. Zeng, J. Keane, and A. Iosifidis, "Adaptive weighting scheme for automatic time-series data augmentation," 2021, *arXiv:2102.08310*.
- [109] E. A. Lopez-Rojas and S. Axelsson, "Money laundering detection using synthetic data," in *Proc. 27th Annu. Workshop Swedish Artif. Intell. Soc.*, Örebro, Sweden, May 2012, pp. 33–40.
- [110] S. Buschjäger, P.-J. Honysz, and K. Morik, "Randomized outlier detection with trees," *Int. J. Data Sci. Anal.*, vol. 13, no. 2, pp. 91–104, Mar. 2022.
- [111] M. Weber, J. Chen, T. Suzumura, A. Pareja, T. Ma, H. Kanezashi, T. Kaler, C. E. Leiserson, and T. B. Schardl, "Scalable graph learning for anti-money laundering: A first look," 2018, *arXiv:1812.00076*.
- [112] T. Suzumura and H. Kanezashi. (2021). *Anti-Money Laundering Datasets*. [Online]. Available: <http://github.com/IBM/AMLSim/>
- [113] K. Singh and P. Best, "Anti-money laundering: Using data visualization to identify suspicious activity," *Int. J. Accounting Inf. Syst.*, vol. 34, Sep. 2019, Art. no. 100418.
- [114] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, Nov. 2008.
- [115] L. Maaten, "Learning a parametric embedding by preserving local structure," in *Proc. 12th Int. Conf. Artif. Intell. Stat.*, Clearwater Beach, FL, USA, Apr. 2009, pp. 384–391.
- [116] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, Mar. 1951.
- [117] J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, Dec. 2000.
- [118] A. Mead, "Review of the development of multidimensional scaling methods," *J. Roy. Stat. Soc. D, Statistician*, vol. 41, no. 1, pp. 27–39, 1992.
- [119] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013, *arXiv:1312.6114*.
- [120] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Proc. 28th Conf. Neural Inf. Process. Syst.*, Montreal, QC, Canada, Dec. 2015, pp. 3483–3491.
- [121] M. Mirza and S. Osindero, "Conditional generative adversarial nets," 2014, *arXiv:1411.1784*.
- [122] E. Brophy, Z. Wang, Q. She, and T. Ward, "Generative adversarial networks in time series: A survey and taxonomy," 2021, *arXiv:2107.11098*.

- [123] J. Yoon, D. Jarrett, and M. Schaar, "Time-series generative adversarial networks," in *Proc. 33rd Conf. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2019, pp. 5508–5518.
- [124] X. Zeng, W. Ouyang, J. Yan, H. Li, T. Xiao, K. Wang, Y. Liu, Y. Zhou, B. Yang, Z. Wang, H. Zhou, and X. Wang, "Crafting GBD-Net for object detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 9, pp. 2109–2123, Sep. 2018.
- [125] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Salt Lake City, UT, USA, Jun. 2018, pp. 7132–7141.
- [126] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever. (2018). *Improving Language Understanding With Unsupervised Learning*. [Online]. Available: <https://openai.com/blog/language-unsupervised/>
- [127] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," 2018, *arXiv:1810.04805*.
- [128] A. Tsantekidis, N. Passalis, A. Tefas, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Forecasting stock prices from the limit order book using convolutional neural networks," in *Proc. IEEE 19th Conf. Bus. Informat. (CBI)*, Thessaloniki, Greece, Jul. 2017, pp. 7–12.
- [129] Z. Zhang, S. Zohren, and S. Roberts, "DeepLOB: Deep convolutional neural networks for limit order books," *IEEE Trans. Signal Process.*, vol. 67, no. 11, pp. 3001–3012, Jun. 2019.
- [130] A. Ntakaris, G. Mirone, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Feature engineering for mid-price prediction with deep learning," *IEEE Access*, vol. 7, pp. 82390–82412, 2019, doi: [10.1109/ACCESS.2019.2924353](https://doi.org/10.1109/ACCESS.2019.2924353).
- [131] N. Woodruff, A. Enshaehi, and B. A. S. Hasan, "Fully-automatic pipeline for document signature analysis to detect money laundering activities," 2021, *arXiv:2107.14091*.
- [132] R. B. D. N. Junior, L. F. Verçosa, D. Macêdo, B. L. D. Bezerra, and C. Zanchettin, "A fast fully octave convolutional neural network for document image segmentation," 2020, *arXiv:2004.01317*.
- [133] J. Han, Y. Huang, S. Liu, and K. Towey, "Artificial intelligence for anti-money laundering: A review and extension," *Digit. Finance*, vol. 2, nos. 3–4, pp. 211–239, Dec. 2020.
- [134] V. Nair and G. E. Hinton, "Rectified linear units improve restricted Boltzmann machines," in *Proc. 27th Int. Conf. Mach. Learn.*, Haifa, Israel, Jun. 2010, pp. 807–814.
- [135] B. Xu, N. Wang, T. Chen, and M. Li, "Empirical evaluation of rectified activations in convolutional network," 2015, *arXiv:1505.00853*.
- [136] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and accurate deep network learning by exponential linear units (ELUs)," 2015, *arXiv:1511.07289*.
- [137] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, "Self-normalizing neural networks," in *Proc. 31st Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Dec. 2017, pp. 972–981.
- [138] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. 13th Int. Conf. Artif. Intell. Statist.*, Sardinia, Italy, May 2010, pp. 249–256.
- [139] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Las Vegas, NV, USA, Jun. 2016, pp. 770–777.
- [140] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller, "Efficient backprop," in *Neural Networks: Tricks of the Trade (Lecture Notes in Computer Science)*, G. Montavon, G. B. Orr, and K.-R. Müller, Eds. Berlin, Germany: Springer, 2012, pp. 9–48.
- [141] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. 32nd Int. Conf. Mach. Learn.*, Lille, France, Jul. 2015, pp. 448–456.
- [142] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *Proc. 30th Int. Conf. Mach. Learn.*, Atlanta, GA, USA, Jun. 2013, pp. 1310–1318.
- [143] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, *arXiv:1607.06450*.
- [144] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Santiago, Chile, Dec. 2015, pp. 1026–1034.
- [145] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Honolulu, HI, USA, Jul. 2017, pp. 2261–2269.
- [146] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov, "Improving neural networks by preventing co-adaptation of feature detectors," 2012, *arXiv:1207.0580*.
- [147] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, Jun. 2014.
- [148] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2006. [Online]. Available: <https://www.deeplearningbook.org/>
- [149] N. Srebro and A. Shraibman, "Rank, trace-norm and max-norm," in *Proc. 18th Conf. Learn. Theory*, Bertinoro, Italy, Jun. 2015, pp. 545–560.
- [150] K. Weiss, T. M. Khoshgofaar, and D. Wang, "A survey of transfer learning," *J. Big Data*, vol. 3, May 2016, Art. no. 9.
- [151] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Proc. 27th Adv. Neural Inf. Process. Syst.*, Montreal, QC, Canada, Dec. 2014, pp. 3320–3328.
- [152] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by back-propagation," in *Proc. 32nd Int. Conf. Mach. Learn.*, Lille, France, Jun. 2015, pp. 1180–1189.
- [153] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko, "Simultaneous deep transfer across domains and tasks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Santiago, Chile, Dec. 2015, pp. 4068–4076.
- [154] L. Hedegaard, O. A. Sheikh-Omar, and A. Iosifidis, "Supervised domain adaptation: A graph embedding perspective and a rectified experimental protocol," *IEEE Trans. Image Process.*, vol. 30, pp. 8619–8631, 2021.
- [155] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [156] H. Sak, A. Senior, and F. Beaufays, "Long short-term memory based recurrent neural network architectures for large vocabulary speech recognition," 2014, *arXiv:1402.1128*.
- [157] W. Zaremba, I. Sutskever, and O. Vinyals, "Recurrent neural network regularization," 2014, *arXiv:1409.2329*.
- [158] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," 2014, *arXiv:1406.1078*.
- [159] N. Passalis, A. Tefas, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Temporal bag-of-features learning for predicting mid price movements using high frequency limit order book data," *IEEE Trans. Emerg. Topics Comput. Intell.*, vol. 4, no. 6, pp. 774–785, Dec. 2020.
- [160] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," 2014, *arXiv:1409.0473*.
- [161] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Trans. Signal Process.*, vol. 45, no. 11, pp. 2673–2681, Nov. 1997.
- [162] M.-T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," 2015, *arXiv:1508.04025*.
- [163] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. 30th Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Dec. 2017, pp. 5998–6008.
- [164] D. T. Tran, A. Iosifidis, J. Kannianen, and M. Gabbouj, "Temporal attention-augmented bilinear network for financial time-series data analysis," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 5, pp. 1407–1418, May 2019.
- [165] O. Kuiper, M. van den Berg, J. van der Burgt, and S. Leijnen, "Exploring explainable AI in the financial sector: Perspectives of banks and supervisory authorities," 2021, *arXiv:2111.02244*.
- [166] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Comput. Surv.*, vol. 54, no. 6, pp. 1–35, Jul. 2021.
- [167] C. Louizos, K. Swersky, Y. Li, M. Welling, and R. Zemel, "The variational fair autoencoder," 2015, *arXiv:1511.00830*.
- [168] A. Gretton, K. M. Borgwardt, M. Rasch, B. Scholkopf, and A. J. Smola, "A kernel method for the two-sample-problem," in *Proc. 19th Conf. Neural Inf. Process. Syst.*, Vancouver, BC, Canada, Dec. 2006, pp. 513–520.
- [169] B. H. Zhang, B. Lemoine, and M. Mitchell, "Mitigating unwanted biases with adversarial learning," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc.*, New Orleans, LA, USA, Dec. 2018, pp. 335–340.
- [170] M. Du, N. Liu, and X. Hu, "Techniques for interpretable machine learning," *Commun. ACM*, vol. 63, no. 1, pp. 68–77, Jan. 2020.

- [171] M. T. Ribeiro, S. Singh, and C. Guestrin, “‘Why should I trust you?’: Explaining the predictions of any classifier,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, San Francisco, CA, USA, Aug. 2016, pp. 1135–1144.
- [172] K. Simonyan, A. Vedaldi, and A. Zisserman, “Deep inside convolutional networks: Visualising image classification models and saliency maps,” 2013, *arXiv:1312.6034*.
- [173] A. Goldstein, A. Kapelner, J. Bleich, and E. Pitkin, “Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation,” *J. Comput. Graph. Statist.*, vol. 24, no. 1, pp. 44–65, Mar. 2015.
- [174] J. H. Friedman, “Greedy function approximation: A gradient boosting machine,” *Ann. Statist.*, vol. 29, no. 5, pp. 1189–1232, Oct. 2001.
- [175] D. W. Apley and J. Zhu, “Visualizing the effects of predictor variables in black box supervised learning models,” 2016, *arXiv:1612.08468*.
- [176] A. Akula, S. Wang, and S.-C. Zhu, “CoCoX: Generating conceptual and counterfactual explanations via fault-lines,” in *Proc. 34th AAAI Conf. on Artif. Intell.*, New York, NY, USA, Feb. 2020, pp. 2594–2601, doi: 10.1609/aaai.v34i03.5643.
- [177] F. Cheng, Y. Ming, and H. Qu, “DECE: Decision explorer with counterfactual explanations for machine learning models,” *IEEE Trans. Vis. Comput. Graph.*, vol. 27, no. 2, pp. 1438–1447, Feb. 2021.
- [178] O. Gomez, S. Holter, J. Yuan, and E. Bertini, “ViCE: Visual counterfactual explanations for machine learning models,” in *Proc. 25th Int. Conf. Intell. User Interfaces*, Cagliari, Italy, Mar. 2020, pp. 531–535.
- [179] F. Hvilshøj, A. Iosifidis, and I. Assent, “ECINN: Efficient counterfactuals from invertible neural networks,” 2021, *arXiv:2103.13701*.
- [180] F. Hvilshøj, A. Iosifidis, and I. Assent, “On quantitative evaluations of counterfactuals,” 2021, *arXiv:2111.00177*.
- [181] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Proc. Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Dec. 2017, pp. 4765–4774.
- [182] L. S. Shapley, “A value for n-person games,” in *Contributions to the Theory of Games*, vol. 2, H. W. Kuhn and A. W. Tucker, Eds. Princeton, NJ, USA: Princeton Univ. Press, 1953, pp. 307–317.
- [183] S. M. Lundberg, G. G. Erion, H. Chen, A. J. DeGrave, J. M. Prutkin, B. G. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee, “From local explanations to global understanding with explainable AI for trees,” *Nat. Mach. Intell.*, vol. 2, no. 1, pp. 56–67, Jan. 2020.



RASMUS INGEMANN TUFFVESON JENSEN received the master’s degree in mathematical-economics from Aarhus University, in 2020, where he is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering. He is also with Spar Nord Bank, Denmark. His research interests include computational finance and operational research.



ALEXANDROS IOSIFIDIS (Senior Member, IEEE) is currently a Professor with Aarhus University, Denmark. He leads the Machine Learning and Computational Intelligence Group, Department of Electrical and Computer Engineering, and the machine intelligence research area with the Centre for Digitalisation, Big Data, and Data Analytics (DIGIT), Aarhus University. He is a member of the IEEE Technical Committee on Machine Learning for signal processing. He is the Associate

Editor-in-Chief of the *Neurocomputing* journal, covering the research area of neural networks, and an Associate Editor of *IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS* and *IEEE TRANSACTIONS ON ARTIFICIAL INTELLIGENCE*.

• • •