



HAL
open science

Are AI-based Anti-Money Laundering (AML) Systems Compatible with European Fundamental Rights?

Astrid Bertrand, Winston Maxwell, Xavier Vamparys

► To cite this version:

Astrid Bertrand, Winston Maxwell, Xavier Vamparys. Are AI-based Anti-Money Laundering (AML) Systems Compatible with European Fundamental Rights?. ICML 2020 Law and Machine Learning Workshop, Jul 2020, Vienne, Austria. hal-02884824v2

HAL Id: hal-02884824

<https://hal.science/hal-02884824v2>

Submitted on 9 Jul 2020 (v2), last revised 12 Nov 2020 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Research paper by

**OPERATIONAL
AI ETHICS**



IP PARIS


telecom-paris.fr/en/ai-ethics

Are AI-based Anti-Money Laundering (AML) Systems Compatible with European Fundamental Rights?

Astrid Bertrand¹, Winston Maxwell², and Xavier Vamparys³

► To cite this version:

Astrid Bertrand, Winston Maxwell, Xavier Vamparys. Are AI-based Anti-Money Laundering (AML) Systems Compatible with European Fundamental Rights? 2020.



Submitted on 19 June 2020

¹PhD candidate, Télécom Paris, Institut Polytechnique de Paris

²Director of Law and Digital Technology Studies, Télécom Paris, i3, Institut Polytechnique de Paris

³Manager of AI Ethics, CNP Assurances; visiting researcher, Télécom Paris, Institut Polytechnique de Paris

Are AI-based Anti-Money Laundering (AML) Systems Compatible with European Fundamental Rights?

Astrid Bertrand¹, Winston Maxwell², and Xavier Vamparys³

¹PhD candidate, Télécom Paris, Institut Polytechnique de Paris

²Director of Law and Digital Technology Studies, Télécom Paris, i3, Institut Polytechnique de Paris

³Manager of AI Ethics, CNP Assurances; visiting researcher, Télécom Paris, Institut Polytechnique de Paris

July 8, 2020

Abstract

Anti-money laundering and countering the financing of terrorism (AML) laws require banks to deploy transaction monitoring systems (TMSs) to detect suspicious activity of bank customers and report the activity to law enforcement authorities. Because the monitoring of customer data to detect money laundering interferes with fundamental rights, AML systems must comply with the proportionality test under European fundamental rights law, as most recently expressed by the Court of Justice of the European Union (CJEU) in the *Digital Rights Ireland* and *Tele2 Sverige - Watson* cases. To our knowledge there has been no analysis as to whether AML systems are compliant with the proportionality test as expressed in these latest cases. Understanding how the proportionality test applies to current AML systems is all the more important as banks and regulators consider moving to AI-based tools to detect suspicious transactions. The objective of this paper is twofold: to study whether current AML systems are compliant with the proportionality test, and to study whether a move towards AI in AML systems could exacerbate the proportionality problems. Where possible, we suggest possible cures to the proportionality problems identified.

Contents

1	AI can help AML become more effective	3
1.1	Money launderers have a high chance of succeeding	3
1.2	Legal barriers to adopting AI-based AML tools	3
2	The proportionality test	4
2.1	The purpose and origin of the proportionality test	4
2.2	Simplifying assumptions	4
2.3	The three steps of the proportionality test	5
2.4	The CJEU <i>Digital Rights Ireland</i> and <i>Tele2 Sverige - Watson</i> cases	5
2.5	The CJEU <i>Ministerio Fiscal</i> case	6
2.6	The Netherlands social security fraud case	6
3	Description of AML measures and the impact of AI	7
3.1	Regulations	7
3.2	The AML process	8
3.3	The shortcomings of current low-tech AML frameworks	10
3.4	How AI is transforming AML frameworks	10
4	Applying each proportionality test to AML transaction monitoring	12
4.1	Do AML systems interfere with fundamental rights?	12
4.2	"Provided by law"	12
4.3	Pursuing a legitimate objective	13
4.4	"Necessary in a democratic society"	13
	"Genuinely effective" ▪ "Less intrusive means" ▪ "Fair balance" ▪ "Adequate safeguards" ▪ Transparency ▪ Explainability ▪ Institutional oversight	
5	Can these proportionality problems be cured?	16
5.1	Problem 1: the law is not specific enough.	16
5.2	Problem 2: you cannot measure effectiveness.	17
5.3	Problem 3: lack of "fair balance" because AML processes conduct "general and indiscriminate" analysis of all transaction data.	17
5.4	Problem 4: lack of transparency and explainability.	17
5.5	Problem 5: the absence of an independent supervisory authority with a 360° vision of the system.	18
6	Conclusion	19

1 AI can help AML become more effective

1.1 Money launderers have a high chance of succeeding

Europol estimates that approximately €200 billion in criminal funds circulate every year in Europe. Yet only 1% of criminal funds are confiscated (Europol [2017]). Anti-money laundering and countering the financing of terrorism (AML or AML-CFT) measures are intended to make it harder for criminals to use traditional financial networks to transfer or invest funds from criminal transactions, but so far results are far from satisfactory. Banks have invested in costly transaction monitoring systems (TMSs) that generate alerts and help bank employees report unusual or suspicious activity to financial crime law enforcement authorities, so-called Financial Intelligence Units (FIUs). But current detection systems generate 90% of false positives, which must then be reviewed by bank employees (IBM [2019]). Each major bank has hundreds of bank employees (or consultants) who have to review thousands of alerts a day to try to separate false positives from real suspicious activities. After this selection is done, suspicious activity reports (SARs) are sent to FIUs, which then decide whether or not to alert the prosecutor's office or the tax, social or customs administration. But FIUs typically review only 10% of the SARs they receive (Europol [2017]), and only a portion of those are forwarded to other law enforcement agencies for action (TRACFIN [2018]).

According to Rob Wainwright, former executive director of Europol: "We have created a whole ton of regulations ... the banks are spending \$20 billion a year to run the compliance regime ... and we are seizing 1 percent of criminal assets every year in Europe" (Paravicini [2018]). AML detection techniques need to improve. Machine learning can help in five ways : automate data collection, enhance client risk scoring and alert prioritization, leverage link analysis, improve segmentation and sharpen anomaly detection. But several obstacles stand in the way, among them fundamental rights and in particular the proportionality test. This paper will examine how the proportionality test imposed by the Charter of Fundamental Rights of the European Union (Charter) and the European Convention on Human Rights (ECHR) applies to existing AML and how the addition of AI-based models would change the proportionality analysis, if at all.

1.2 Legal barriers to adopting AI-based AML tools

The low implementation of AI techniques in banks today is attributable to several factors, including technical and organizational barriers¹. This paper focuses on one major legal barrier², which is whether AI-based systems would comply with the proportionality test of the CJEU. The proportionality test applies to any government-imposed system designed to fight crime that also interferes with fundamental rights such as privacy. In order to determine whether new AI-based AML systems could satisfy the proportionality test, we first need to ask whether *current* AML systems satisfy the proportionality test, particularly as that test was expressed in two CJEU cases involving the processing of customer data by telecommunications operators. Those cases imposed limits on what governments can ask private entities to do with customer data in order to help law enforcement agencies. Our analysis shows that current AML systems probably fail the proportionality test as expressed in the CJEU's *Digital Rights Ireland* and *Tele2 Sverige - Watson* cases. Adding AI will exacerbate the existing problems, and create a new one in the form of explainability. The

¹These barriers include the existence of complex legacy systems, the non-availability of appropriate training data, and the existence of organizational silos within large banks.

²Other legal barriers include the current structure of AML regulations, which divide detection and enforcement responsibilities between banks and government authorities, as well as the emphasis in current AML regulations on rule-based scenarios and human review.

defects can likely be cured through improved legislation calling for enhanced institutional supervision and transparency. However, a major defect is the inability to measure the effectiveness of various AML systems, whether or not they use AI. This defect must be cured before AI can be introduced.

Our focus on proportionality for AML will shed light on many other use cases involving AI to fight crime: fraud detection, predictive policing, airline passenger screening, cyber-security, and counter-terrorism to name a few. The proportionality test will apply to all of them because they can adversely impact fundamental rights, particularly the right to privacy. Consequently, the contribution of this paper extends well beyond the AML use case.

The remainder of the paper is structured as follows: section 2 provides an overview of the proportionality test; section 3 gives an overview of AML regulation and processes; section 4 applies the proportionality test to AML systems; section 5 provides suggested cures for the identified problems; section 6 concludes.

2 The proportionality test

2.1 The purpose and origin of the proportionality test

Measures designed to fight crime must be analyzed under the proportionality test to ensure that government powers, for example powers to conduct electronic surveillance, do not unduly restrict privacy or other fundamental rights such as non-discrimination, freedom of expression, or a right to a fair trial ([Hickman \[2008\]](#)). The proportionality test is a meta-principle that flows from national constitutions, the EU treaties, the Charter, and ECHR ([Sauvé \[2017\]](#)). The principle trickles down into the various laws and regulations that are involved in AML, including the GDPR³, the Police-Justice Directive⁴, and the AML Directives⁵. Consequently, when a bank implements a TMS, the bank will have to conduct a data protection impact assessment under article 35 of the GDPR to evaluate "the necessity and proportionality of the processing operations in relation to the purposes"⁶. When a Member State such as France enacts a law or regulation providing for the processing of AML-related data by the French FIU Tracfin, the processing must be "necessary and proportionate" under the Police-Justice Directive⁷. In any given situation, proportionality will be both a requirement flowing from fundamental rights texts, and a requirement flowing from a particular law or directive, such as the GDPR or the Police-Justice Directive.

2.2 Simplifying assumptions

The multi-layered source of proportionality can get confusing, which is why we propose to simplify matters by looking at the proportionality test as it flows from fundamental rights texts such as the ECHR and the Charter. The proportionality test is expressed in slightly different ways by different courts, constitutions and treaties. To simplify the discussion below, we will assume that there is a

³Regulation (EU) 2016/679 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation or GDPR).

⁴Directive (EU) 2016/680 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and on the free movement of such data.

⁵See note 3.1 above

⁶Article 35(7)(b), GDPR

⁷Article 4(2)(b) Police-Justice Directive

single proportionality test, which is not technically correct. Different courts take slightly different approaches (Hart [2015]). But the differences are not important enough to affect the conclusions made in this paper.

2.3 The three steps of the proportionality test

The proportionality test can be broken down into three tests that must be cumulatively satisfied. The first test consists in asking whether the measure has been provided for in a law or regulation that is precise, understandable and has been adopted pursuant to democratic processes⁸. The second test consists in asking whether the measure pursues a legitimate objective such as a fundamental right or another pressing social need. The third test consists in asking whether the measure is "necessary in a democratic society." This is the most tricky part of the proportionality review. It involves several sub-tests, such as whether the measure is genuinely effective and is the least intrusive measure available. Other sub-tests involve asking whether there is a fair balance between the rights at stake, and whether the safeguards, including institutional oversight and transparency, are sufficient to mitigate any adverse impacts on fundamental rights.

The three tests (and related sub-tests) are summarized in figure 1.




Figure 1. The proportionality test can be broken down into three tests. To satisfy the proportionality test, the answers to all the questions in the graph must be yes.

2.4 The CJEU *Digital Rights Ireland* and *Tele2 Sverige - Watson* cases

Several cases shed light on how the proportionality test would apply in practice to AML systems. In 2014, the Court of Justice of the European Union (CJEU) annulled a 2006 directive requiring Member States to ensure that telecommunications operators retain traffic and location data for up to two years in order to assist law enforcement authorities and intelligence agencies investigate crimes or prevent terrorism.⁹ The CJEU found that the level of interference with privacy rights was particularly high, because the directive required operators to store all traffic and location data of users, regardless of whether the users posed a particular risk, or whether the data had a particular link to an investigation. It also failed to provide limits on the persons who could access the data. Because of its generality and lack of adequate safeguards, the directive went beyond what was strictly necessary and failed the proportionality test. Two years later, the CJEU declared illegal two

⁸Ahmet Yildirim v. Turkey, ECtHR n° 31111/10, December 18, 2012.

⁹*Digital Rights Ireland v. Minister for Communications*, CJEU, Joined Cases C-293/12 & C-594/12, (Apr. 8, 2014)

national laws containing similar provisions requiring telecom operators to store all data in case it is needed by law enforcement or national security agencies¹⁰. The CJEU found that the general and indiscriminate retention of all traffic and location data exceeds what is strictly necessary. In order to be acceptable, *“the national legislation must be based on objective evidence which makes it possible to identify a public whose data is likely to reveal a link, at least an indirect one, with serious criminal offences, and to contribute in one way or another to fighting serious crime or to preventing a serious risk to public security. Such limits may be set by using a geographical criterion where the competent national authorities consider, on the basis of objective evidence, that there exists, in one or more geographical areas, a high risk of preparation for or commission of such offences”*¹¹. This wording is extremely important for AML, because it suggests that wholesale analysis of all customer transaction data would be excessive, and that there would have to be some objective link between the population whose data is being analyzed and money-laundering risks. The CJEU also said that there must be safeguards, such as informing the relevant individuals as soon as it is possible to do so without jeopardizing an ongoing criminal investigation. The system would also require institutional oversight.

2.5 The CJEU *Ministerio Fiscal* case

Another recent CJEU case¹² involved a Spanish law that required operators to collect and store identification data relating to persons who purchase SIM cards, and make the identification data available to law enforcement authorities if so ordered. The case involved a relatively minor offence, theft of a mobile phone. The CJEU had to determine whether the Spanish law preserved a "fair balance" under the proportionality test insofar as the law did not only apply to serious crimes, but to more minor offences such as theft of a mobile phone. The CJEU found that the "fair balance" sub-test was respected in this case because the data involved, only the name and address of the customer, was much less intrusive than the detailed traffic and location data involved in the *Digital Rights Ireland* and *Tele2 Sverige-Watson* cases. According to the CJEU, the more detailed and intrusive the data, the more serious must be the crime: *“In accordance with the principle of proportionality, serious interference can be justified, in areas of prevention, investigation, detection and prosecution of criminal offences, only by the objective of fighting crime which must also be defined as ‘serious’. By contrast, when the interference that such access entails is not serious, that access is capable of being justified by the objective of preventing, investigating, detecting and prosecuting ‘criminal offences’ generally.”*¹³. When applied to AML, this means that the analysis of detailed transaction data which entails a high level of interference with individual rights would only be permitted for detection of *serious* crimes.

2.6 The Netherlands social security fraud case

The proportionality test was applied recently by a court in the Hague¹⁴, Netherlands, with regard to an algorithm used by the social security authorities to detect potential social security fraud. Netherlands law authorizes the government to collect data from various government databases in order to create risk profiles showing the likelihood that a given individual is cheating on social

¹⁰ *Tele2 Sverige v. Post- och telestyrelsen*, and *Secretary of State v. Watson*, CJEU, Joined Cases C-203/15 & C-698/15, (Dec. 21, 2016)

¹¹ *Id.* point 111.

¹² CJEU, Case C-207/16, *Ministerio Fiscal*, 2 October 2018.

¹³ CJEU, Case C-207/16, *Ministerio Fiscal*, 2 October 2018, points 56 and 57.

¹⁴ Rb. Den Haag 5 February 2020 (Nederlands Juristen Comité voor de Mensenrechten/Staat Der Nederlanden) (Neth.), <https://uitspraken.rechtspraak.nl/inziendocument?id=ECLI:NL:RBDHA:2020:865>

security benefits. The District Court of the Hague found that the measure was “provided for by law” because the law was sufficiently detailed, and that the objective of fighting welfare fraud was a compelling social interest, satisfying the first and second test. When it applied the “necessary in a democratic society” test, the court found that the system was surrounded by inadequate safeguards and in particular that it lacked transparency. First, individuals were not informed of the existence of the system, nor given a general understanding of how data about them was being used. Second, neither the court nor individuals targeted by algorithmic risk profiles were able to understand how the proprietary model operated and how it reached a particular score. This kind of understanding (often referred to as “explainability”) is necessary to permit individuals to defend themselves, and to permit courts to verify the presence or absence of discrimination. Because the system lacked this transparency, it failed the proportionality test. The risk score reports did not lead to automatic decisions. The reports were reviewed by employees who then decided whether to investigate. The Netherlands court did not examine the case under the GDPR or the Police-Justice Directive, but relied solely on the proportionality text of the European Convention of Human Rights. The parallels between the Netherlands social security fraud case and AML measures are striking, because one of the functions of TMSs is to build risk profiles, similar to the fraud risk profiles created by the Netherlands social security authorities. In both situations, human reviewers evaluate the risk scores or alerts generated by the system, and decide whether to take further action.

3 Description of AML measures and the impact of AI

3.1 Regulations

In Europe, AML has given rise to six directives¹⁵ and countless national laws and regulations. International cooperation is ensured by the Financial Action Task Force (FATF). AML legislation imposes several tasks on banks: performing KYC which includes identifying any politically exposed person (PEP); using a TMS to filter transactions (TRACFIN [2018], Dupont et al. [2020]); detecting transactions related to a high-risk country, and detecting and reporting incoherent or anomalous transactions based on the organization’s knowledge of the customer. Banks must also implement processes to ensure that no payments are made to persons or entities subject to international sanctions (Dupont et al. [2020]).

The originality of AML is that banks and other private entities are required to actively look for suspicious activity and report it to FIUs without informing their customers. Reporting entities, and particularly banks, credit and payment institutions and insurance companies play a key role in the fight against money laundering and terrorism financing, as law enforcement authorities heavily depend on the information provided by such entities. They are often seen as significant “partners” of such authorities in the global AML system. The French FIU called its cooperation with reporting entities a “genuine and comprehensive “public-private” partnership”. From the banks’ standpoint, the “partnership” is seen more as a one-way street, banks being imposed law-enforcement tasks that would more naturally fall on the police. The cost of compliance with AML regulations is significant, with no obvious upside for financial institutions, particularly since the latter have a tradition of bank secrecy vis à vis their customers. For this reason, financial institutions generally refuse to see themselves as “advanced substitutes for law enforcement authorities and justice” (Lascoumes

¹⁵First directive 91/308/CE of 28 June 1991, second directive 2001/97/CE of 4 December 2001, third directive 2005/60/CE of 26 October 2005, fourth directive 2015/849/CE of 15 May 2015, fifth directive (UE) n°2018/843 of 30 May 2018 and sixth directive (UE) n° 2018/1673 of 23 October 2018 on the prevention of the use of the financial system for the purpose of money laundering and terrorist financing.

and Favarel-garrigues [2006]) for AML purposes. In addition, financial institutions receive little feedback on the actions taken after the SAR was sent to a FIU, which prevents institutions from implementing more efficient and less expensive AML processes.

The AML system changed radically with the adoption of the third and fourth European directives on AML. Prior directives imposed specific predefined rules that reporting entities were obliged to follow, without leaving the freedom to modify or adapt these rules to the specific situation of each such entity or the circumstances of a particular transaction. Such rules were easy to apply, in particular by compliance officers whose job was limited to checking that all required boxes had been ticked, but the system proved to be inefficient.

The third Directive on AML introduced a more flexible “risk-based” approach amplified by the fourth Directive. Such Directive provides that the *“risk of money laundering and terrorist financing is not the same in every case. Accordingly, a holistic, risk-based approach should be used. The risk-based approach is not an unduly permissive option for (...) obliged entities. It involves the use of evidence-based decision-making in order to target the risks of money laundering and terrorist financing facing the Union and those operating within it more effectively”*¹⁶. The most relevant risk factors for a bank are its customers, countries of operation, types of products or services and distribution channels. Accordingly, AML frameworks of reporting entities operating in numerous countries and selling complex products through various intermediaries must be more complex – and cost intensive – than those of reporting entities directly selling plain-vanilla products. Each reporting entity must determine the appropriate risk profile of each of its customer or transaction and apply the relevant level of scrutiny (simplified, standard, reinforced or complementary)¹⁷.

3.2 The AML process

Figure 2 explains the usual frameworks deployed by banks to respond to their AML obligations¹⁸.

Every incoming transaction is first screened by an automated review system that uses the information gathered on the client profile, the transaction details, information concerning watch-listed countries and other motives of alerts. Other data sources such as market activities, trade-based data or even social media and news feed are sometimes employed, but their usage is highly dependent on the bank’s ability to exploit them. This automated review system will decide to generate an alert or not for a transaction, either based on deterministic rules of type if-then or otherwise relying upon more sophisticated AI models. All the generated alerts then follow a usually three - maybe two - steps review by compliance experts who choose whether to escalate the transaction or to close it. Alerts that go through all of these review stages are then consolidated into SARs (suspicious activity reports) that are forwarded to the Financial Investigation Units for final investigation to determine the actual threats.

Automated Review Systems usually perform several upstream tasks to decide whether or not to generate an alert for a particular transaction. The first task consists in dividing customers into specific segments based on common behavioral attributes. This process is usually driven by specific industrial knowledge and by rules and thresholds applied to the industry of the clients, their type of

¹⁶Recital 22 of the Directive (EU) 2015/849 of 20 May 2015 on the prevention of the use of the financial system for the purposes of money laundering or terrorist financing.

¹⁷Art. 561 and seq. of the French Monetary and Financial Code. On the requirements of transaction monitoring, see New York Department of Financial Services Regulations Part 504 Transaction Monitoring and Filtering Program Requirements and Certifications.

¹⁸Icons were taken from <https://icons8.com/>

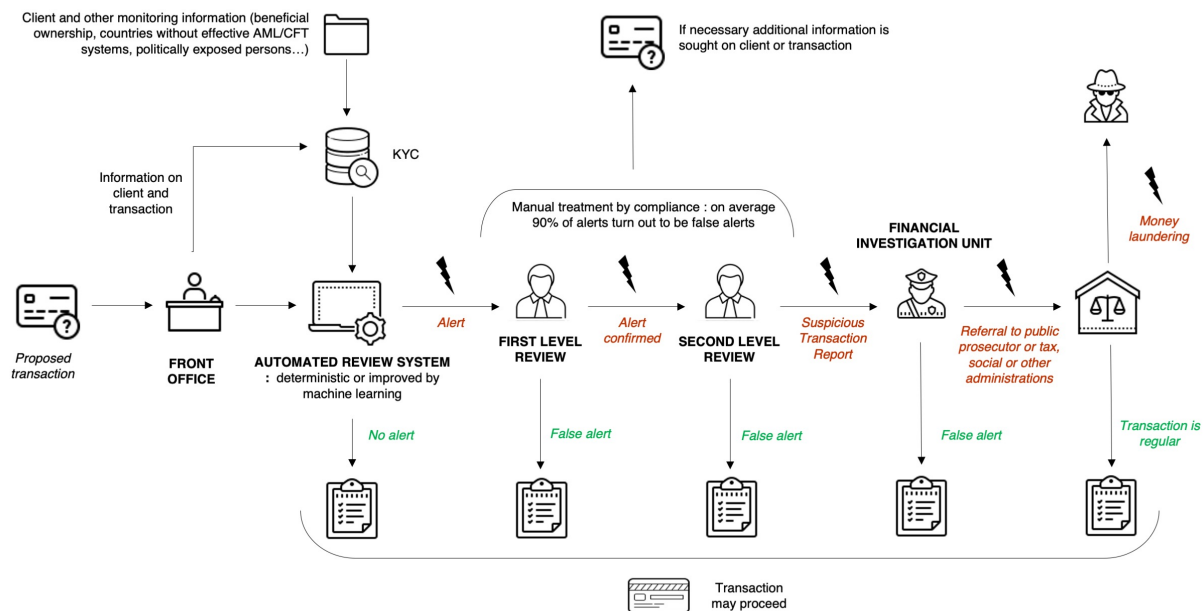


Figure 2. Graph representing the AML processes

business, size, etc... After that, the review system will execute an anomaly detection to see if the incoming transaction falls outside the “normality” scenario for that customer considering his segment profile. This is usually set up by creating thresholds that represent a certain level of deviation from the typical client behavior in the segment. The following step is alert prioritization. As the risk-based approach advocated by regulators became more popular among banks, many models were developed to rank alerts according to the likelihood that they actually correspond to a real risk. This is usually done by calculating a risk-score based on static rules, for instance regarding the transaction amount. Lastly, visualisation tools gathering all useful data for the review can be leveraged, building for example upon graph structures. However, these methodologies are very recent and this last step, or lack thereof, actually constitutes a real issue in standard AML systems where information concerning a specific alert is fragmented, incomplete, or inconsistent. Figure 3 presents the different tasks embedded in the Automated Review System to generate alerts¹⁹.

¹⁹Icons were taken from <https://thenounproject.com/>

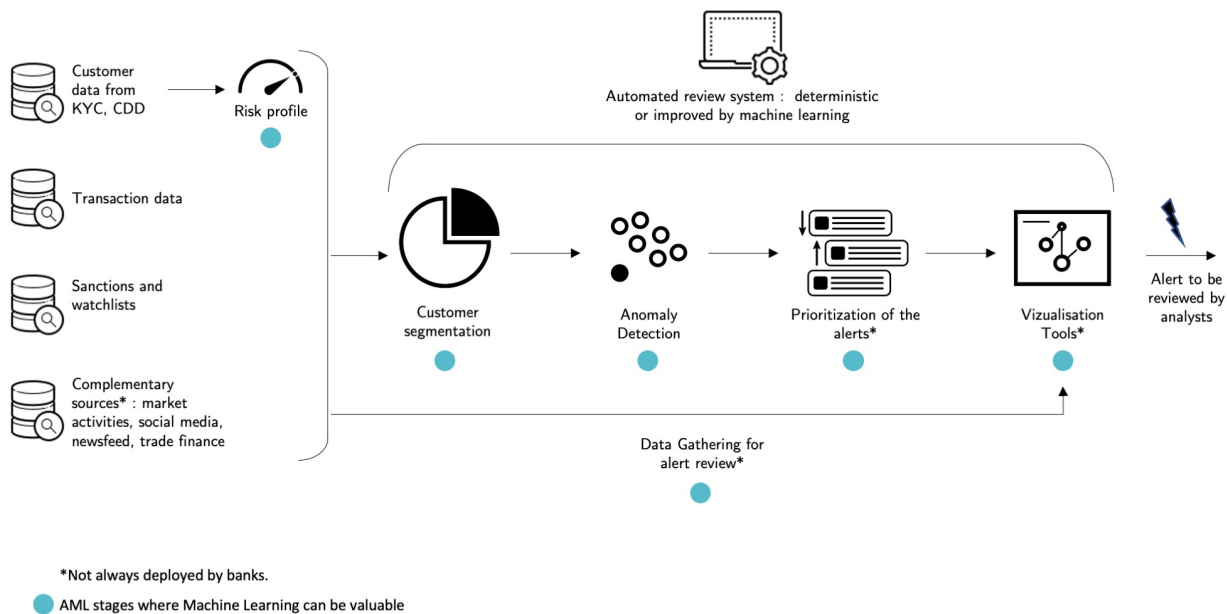


Figure 3. Graph representing the Automated Review System

3.3 The shortcomings of current low-tech AML frameworks

Although machine learning is gradually being adopted by banks, simple rule-based models and manual filtering of alerts still prevail today. It is estimated that about 50% of banks still use manual methods to comply with anti-money laundering requirements ([ACAMS - LEXIS NEXIS \[2015\]](#)). Yet, low-tech, rules-based AML mechanisms exhibit significant limitations. Rules are mainly deterministic, static, and difficult to maintain manually. As a result, AML systems are congested with redundant and obsolete rules which lead to out-of-place alerts, to a huge and unnecessary workload for compliance teams and to a poor effectiveness of AML measures. To cope with the increasing number of alerts within the constraints of human analysis, financial institutions are hiring more and more compliance officers, which significantly increases costs.

3.4 How AI is transforming AML frameworks

In recent years, we have seen a growing interest in AI to assist analysts in highly repetitive AML compliance tasks on the one hand and improve the performance of AML frameworks on the other hand. As shown in Figure 3, Machine Learning can be introduced in several parts of the TMS to improve the efficiency of each task. More specifically, we can distinguish five ways in which AI is transforming current AML systems: automate data collection, enhance the client risk scoring and the alert prioritization processes, leverage link analysis, improve segmentation and sharpen anomaly detection either by identifying known suspicious patterns or by discovering new ML patterns.

Automate data gathering : Machine Learning techniques like Natural Language Processing (NLP) and Optical Character Recognition (OCR) are being progressively appropriated by financial institutions to tap into external unstructured data and enrich the understanding of investigators. Such techniques are very helpful to provide the context and sentiment of a newspaper article or blog post, and other public sources content etc.

Fine-tune Alert prioritization and Client Risk scoring: Using supervised algorithms like decision trees, random Forests and logistic regressions, etc... the processes of customer risk-scoring and alert prioritization can be fine-tuned. For instance, ([Wang and Yang \[2007\]](#)) developed a money-laundering risk evaluation tool using decision trees.

Leverage Link analysis: Social Network Analysis is an emerging field in AML useful both for improving the visualisation of information during investigations and for automatically analyzing inferences between parties. This approach is very powerful in the fight against money laundering as it captures the essence of money laundering schemes i.e. cash flow relationships. For instance, ([Weber et al. \[2019\]](#)) demonstrates the interest of graph structures for capturing relational information by establishing the efficiency of a combination of graph convolutional networks and random forests to recognize illicit Bitcoin transactions. ([Savage et al. \[2016\]](#)) also leverages graph structures to uncover illicit "communities" according to their relationships with known compromised actors in the network.

Improve segmentation: Unsupervised learning can be particularly useful to drive more "intelligent segmentation" by helping compliance experts detect behavioral patterns in the data otherwise invisible during manual review. Clustering techniques like K-means algorithm and PCA are unsupervised machine learning techniques that appear in the AML literature. These AI-based techniques are also promising for partially automating the segmentation model and making it more resistant to variations in the data.

Improve anomaly detection:

- **Detection of known suspicious patterns**

Through supervised learning, models can be trained to distinguish suspicious from normal transactions using previous sanctions cases. These algorithms can be very useful to enhance the anomaly detection phase shown in Figure 3. An example of such models can be found in ([Jun Tang and Jian Yin \[2005\]](#)), where the author designed a supervised classifier to sort out unusual transactions, choosing a Support Vector Machine (SVM) because of its ability to function among high dimensionality heterogeneous data sets. However, the dependency of supervised algorithm on training data rises important challenges for their use in AML. The first is the absence of ground truth on what was in fact a suspicious transaction. The bank will only know what it previously labelled as suspicious, not what actually turned out to be suspicious after investigation by the FIU, or what was entirely missed. The second challenge is the imbalance in training examples, the proportion of suspicious transactions representing only a tiny fraction of the volume of total transactions in the training data set. Finally, supervised models can only detect patterns that are similar to the identified criminal cases given in the training set, whereas money laundering schemes are constantly and rapidly evolving.

- **Uncover new patterns**

Unsupervised learning brings forth possibilities to discover new money laundering patterns otherwise too complicated for humans to notice. Examples of such models in the literature include the works of ([Le Khac and Kechadi \[2010\]](#)), ([Paula et al. \[2016\]](#)), ([Liu et al. \[2011\]](#))

and (Jun Tang and Jian Yin [2005]). (Le Khac and Kechadi [2010] designed a K-means based clustering technique to distinguish suspicious from non-suspicious cases. (Paula et al. [2016]) applied a Deep learning AutoEncoder to detect exporting corporations displaying signs of anomalous behavior. (Liu et al. [2011]) used an unsupervised clustering algorithm designed as a combination of BIRCH and K-means to identify abnormal transactions. In complement to his supervised SVM algorithm, (Jun Tang and Jian Yin [2005]) also investigates an unsupervised one-class SVM for unusual patterns discovery that features the advantage of not requiring training data on the label to perform classification.

If we take the effectiveness metrics employed by banks, namely the reduction of false positive and the ratio of alerts escalated to SAR on total number of alert, Machine Learning methods have proven themselves to be valuable by reducing the false positive rate by 20 to 30% (Weber et al. [2018]).

4 Applying each proportionality test to AML transaction monitoring

4.1 Do AML systems interfere with fundamental rights?

A threshold question is the level of interference with fundamental rights. AML processes are particularly intrusive: they analyze all transaction data without exception, similar to CJEU's *Tele2 Sverige - Watson* case, they result in risk profiles, similar to the Netherlands social security fraud case. The objective of processing is to detect and report criminal activity, processing that normally requires special safeguards under the GDPR²⁰. The level of intrusiveness is already quite high under existing AML processes, because current rule-based TMS algorithms already analyze transaction data to create risk profiles and alerts. As we saw from the Netherlands social security fraud case, the use of machine learning techniques can exacerbate the problem because of the opacity of the algorithms. AML systems impact on privacy rights, because they involve analysis of massive and particularly sensitive bank transaction data, the creation of individual profiles, and possible reports of suspected criminal offences to law enforcement authorities. They may also interfere with the right to non-discrimination, because TMSs may generate more alerts for certain groups of the population than for others. Finally, AML systems may affect the right to an effective remedy because individuals will not be informed of the processing, and even if they are informed, they may be unable to challenge the system due to its complexity.

4.2 "Provided by law"

The first test, "provided by law", requires a law or regulation describing in reasonable detail the type of tool required, the type of data processed, and the safeguards surrounding the use of the tool, so that citizens and other political stakeholders can understand and react, holding the government accountable. Current TMS requirements are described in general terms by regulations, for example (ACPR & TRACFIN [2018]). But the specificity is left to the banks, who develop complex TMSs to satisfy the regulatory authorities' interpretation of the published regulations. There is considerable leeway in the interpretation, and regulators have a tendency to require state-of-the-art detection systems comparable to what other peer banks have implemented. The state

²⁰Article 10 of the GDPR provides that Processing of personal data relating to offences shall be carried out only under the control of official authority or when the processing is authorised by Union or Member State law providing for appropriate safeguards for the rights and freedoms of data subjects.

of the art will constantly evolve based on industry practice. Consequently there is a risk that systems will become more and more intrusive without any clear limits set by law²¹. The lack of specificity of current law is striking when compared with the regulations on storage of customer data by telecom operators²², or the Netherlands law that authorized the use of algorithms to predict social security fraud, both of which provide a high degree of specificity on the data that may be processed and the safeguards that must be applied. More specific regulation on AML systems would be required not only under the "provided by law" branch of the proportionality test, but also from the standpoint of banks' compliance with the GDPR. Article 10 of the GDPR prohibits processing of data relating to criminal offences by private entities in the absence of a specific law containing appropriate safeguards. Also, one of the criticisms made by data protection authorities of current AML systems is that regulators push banks to go beyond what is strictly required by law, a phenomenon known as "gold plating" ([Article 29 Working Party \[2011\]](#)). Gold plating is illegal under the GDPR, because banks are required to ensure that their processing is not excessive. To make sure banks know what to do, and do not go beyond what is required by law (since going beyond what is necessary would be illegal), the legal requirements need to be specific.

4.3 Pursuing a legitimate objective

The second test, pursuing a legitimate objective and pressing social need, is easy to satisfy for AML measures, because AML is part of the fight against serious crime.

4.4 "Necessary in a democratic society"

The third test, "necessary in a democratic society," is the most difficult.

4.4.1 "Genuinely effective"

A key element of proportionality is whether a given approach is genuinely effective in achieving the desired objective and is the least intrusive means ([EDPS \[2017\]](#)). But how do you measure effectiveness? As we noted above, there is a major problem since banks do not know the real results of their notifications to FIUs. The ground truth remains hidden from banks, so they have to imagine other ways of measuring effectiveness. The most frequent approach is to compare the performance of the tool based on the number of alerts that are subsequently "converted" by bank reviewers into SARs. But this is an imperfect metric because it does not capture whether a given approach really helps law enforcement authorities identify and confiscate criminal funds and prosecute criminals, which is, after all, the objective of the law. This defect exists already for current AML methods. Banks only have a vague notion of their effectiveness in the ultimate detection and prosecution of criminals, because of the absence of detailed feedback from FIUs. According to Europol ([Europol \[2017\]](#)), FIUs investigate only 10 percent of the SARs they receive²³. This suggests that most SARs sent to FIUs are basically useless, but nevertheless involve intrusive processing of personal data. Adding machine learning does not fundamentally change the nature of the problem. In the absence of reliable metrics showing the utility of the SARs for the pursuit of criminals it appears impossible to affirm that current AML methods are "genuinely effective" for purposes of the proportionality test. Adding AI to

²¹ AML laws state that TMS systems must always comply with the GDPR, but the GDPR does not set hard limits. Banks justify their TMS under the "required by law" legal basis of the GDPR. If banks believe that the regulator requires more intrusive processing to be compliant with law, the processing will appear compliant under the GDPR. The reasoning becomes circular.

²² For example French Decree n° 2006-358 of 24 March 2006).

²³ The percentage in France is nearer 20 percent.

AML systems will make the problem even more visible, but will not change the nature of the problem.

4.4.2 "Less intrusive means"

Related to the "genuinely effective" question is the question of whether there are less-intrusive means reasonably available that would achieve the same level of effectiveness while creating a lower impact on fundamental rights (EDPS [2017]). This would require analyzing several scenarios and comparing their relative effectiveness and impact on fundamental rights. The optimal scenario would be the one that maximizes both effectiveness and protection of fundamental rights (Alexy [2014], Hickman [2008], Portuese [2013]). But as noted above, maximizing effectiveness requires some reliable way of measuring effectiveness, which currently does not exist. When it comes to measuring the impact on fundamental rights, qualitative scoring methods can be used that assess impacts based on different criteria (Alexy [2014], J. Maxwell [2017], Grazia Porcedda [2011]):

- the level of sensitivity of the data involved, and in particular whether the data can reveal detailed profiles of people's lives;
- whether the system targets only some data of some customers, or whether the system analyzes all data of all customers;
- who has access to data;
- institutional guarantees;
- the level of transparency and explainability;
- security.

Other fundamental rights would also have to be assessed, including the right to non-discrimination (does the tool create clusters based on the geographic origin of names?), and the right to an effective remedy (the ability to challenge algorithmic outcomes). Each approach would be scored in terms of effectiveness and in terms of fundamental rights protection, and in very rough terms, the approach with the highest aggregate score (effectiveness score plus fundamental rights protection score) would emerge as the approach most likely to satisfy the "least intrusive means" test²⁴. Existing AML measures have been in existence for more than a decade. Asking whether there are less intrusive means available 10 years after the introduction of the measures is probably not politically useful because there is no political appetite to go backwards on AML protections. However, there is a recognition that current AML methods have to evolve to become more effective (Europol [2017]), which will lead to a debate on the introduction of AI-based systems. In the context of this debate, an evaluation of alternative measures, their effectiveness and their respective impacts on individual rights will be necessary.

4.4.3 "Fair balance"

The "fair balance" enquiry consists of ensuring that weights of the respective rights are roughly comparable, and that one right does not completely dominate the other. The fair balance enquiry ensures that the essence of each fundamental right is preserved. In the CJEU's *Ministerio Fiscal*, the court said that intrusive data processing can be justified only for detection of *serious* crimes. The scope of AML has been considerably broadened to cover not only drug trafficking, trafficking in

²⁴This is a huge oversimplification. Among other things, each fundamental right would have to receive a minimum score, ensuring that the "essence" of the relevant right is preserved as well.

human beings, corruption, terrorism (all of which are serious crimes), to any crime with a potential sentence of more than one year, including tax fraud. Potentially VAT fraud by the local bistro could be covered. This raises the question of whether the scope of criminal activity should be more limited, so that bank systems look only for serious crimes. The addition of AI does not affect this analysis. Another important consideration in the "fair balance" enquiry is whether AML automatically fails because of generalized and indiscriminate processing of all customer data. AML systems analyze all transaction data of all customers in a generalized and indiscriminate fashion, exactly what the CJEU said was illegal in its case law involving telecommunications operators. The *Tele2 Sverige-Watson* case suggests that wholesale processing of all data is by definition disproportionate, and would cause AML to immediately fail the fair balance test.

4.4.4 "Adequate safeguards"

As part of the "necessary in a democratic society" test, courts will always ask whether the measure is surrounded by adequate safeguards. The types of safeguards are generally those specified by the GDPR: data minimisation, limitation on who has access, limiting the storage, security measures, transparency, and accountability. For government-imposed measures, such as police or national security surveillance, courts focus on the existence of independent institutional oversight. Existing AML measures suffer from two weaknesses which are outside the control of banks. The first is transparency.

4.4.5 Transparency

In principle, individuals must be informed that they have been singled out as posing a risk of criminal activity. However, AML legislation prohibits this, for the understandable reason that tipping off an individual may interfere with a criminal investigation. The compromise accepted by courts is that individuals be informed individually as soon as possible once the information no longer poses a threat for the current investigation. This suggests that AML regulations should be modified to provide that banks must inform customers that have been the subject of a SAR a certain time after the SAR has been transmitted to the FIU, unless the FIU specifically says that doing so would interfere with an ongoing investigation. In light of the fact that a majority of SARs are never investigated by FIUs, it seems disproportionate to impose a permanent black-out even for the SARs that are never used.

4.4.6 Explainability

This problem of transparency is independent of the existence of AI, but AI adds a new wrinkle linked to the explainability of algorithms (Beaudouin et al. [2020]). AI can create new clusters of bank customers based on risks and relationships that are not visible to humans. Customers will have a right to understand why they were put in one cluster and not in another, and why certain scenarios were applied to them. As we learned in the Netherlands social security fraud case, individuals have a right to be informed that risk profiles are being created, and they have a right to understand why an algorithm attributed a particular risk score to them and created an alert. This is true even if the algorithm's recommendation is subsequently reviewed and approved by humans²⁵. There are two reasons for transparency: the first is to make sure individuals have an effective right to challenge what the banks and FIUs are doing. To make an effective challenge, individuals need to be able to challenge the tool relied on by officials. The second is to make sure regulators can ensure that the algorithm is operating as intended, for example it is not creating clusters based on protected features like national origin, religion or gender. It is well known now that machine learning algorithms can

²⁵The Netherlands scoring algorithm did not lead to an automatic decision – its recommendations were reviewed by humans.

discriminate even if they are designed not to. Constant verification is needed to make sure algorithms are not learning to discriminate in unpermitted ways. Explainability is a major concern for AI-based algorithms deployed by banks generally (Dupont et al. [2020]) and more specifically for AML systems.

4.4.7 Institutional oversight

The second major safeguard that appears missing for AML is institutional oversight. When courts apply the proportionality test to government surveillance measures, an important factor is whether an independent court or commission reviews the process regularly to assess effectiveness and continued respect for fundamental rights. In the case of AML, the bank's TMS would be subject to oversight by a bank regulatory authority and by a data protection authority, each authority with different priorities. But the oversight of these two regulatory authorities would typically not extend to law enforcement authorities who use the SARs to pursue criminals. The independent institutional oversight therefore remains partial and incomplete, ignoring the law enforcement part of the equation. While some data protection authorities in theory have power to conduct enquiries on the law enforcement side, this power depends heavily on the existence of specific legislation. In many cases, the data protection authority's powers over police and national security processing are quite limited. By contrast, in the field of national security surveillance for example, there typically exist special independent commissions that have as their mission the evaluation of intelligence gathering technologies deployed by the State²⁶. The independent commission is often asked to prepare annual reports to parliament and to approve in advance the use of certain technologies. An important role of the commission is to evaluate the effectiveness of law enforcement technologies. For AML, the institutional oversight problem is independent of the use of AI. However, as is the case in law enforcement and counter-terrorism, the use of AI in AML opens new risks for citizens that are not yet fully understood, which makes institutional oversight all the more necessary.

5 Can these proportionality problems be cured?

The objective of this paper was to apply the proportionality test to AML systems and identify problems. Proposing specific cures to the proportionality problems is beyond the scope of this paper. Nevertheless, we present below several initial thoughts on potential cures.

5.1 Problem 1: the law is not specific enough.

The first problem relates to the "provided by law" test. More specificity is needed in AML laws to describe the systems the banks are required to put into place to detect suspicious transactions. This is true whether or not AI is used. One objection to publishing a detailed regulation is that it would help money launderers understand the system and avoid detection. However this problem can be overcome by describing the overall functionalities of the system, and the data it may use, in a public regulation, and keeping the specific details of the algorithm in a confidential government ruling. This approach is already used for cybersecurity. In France, for example, the cybersecurity requirements for a specific site of critical national importance will be kept in a document classified as a defense secret. The generic cybersecurity requirements for a given category of sites of critical national importance (for example banks) will be specified in a public regulation.

²⁶An example is the French Commission on the Supervision of Intelligence Gathering Techniques, the CNCTR (Commission nationale de Contrôle des Techniques de Renseignement)

5.2 Problem 2: you cannot measure effectiveness.

Curing this problem requires developing a quality metric that would be applied by law enforcement agencies to the SARs they receive from banks, both SARs developed under existing systems and SARs developed with the help of AI tools. The quality metric would indicate the SAR's actual utility in confiscating criminal funds and prosecuting money launderers. The proportionality test looks at the social utility of the measure, for example the seriousness of the crime. The SAR quality metric should therefore capture two dimensions: the seriousness of the crime on the one hand, and the relative utility of the SAR for stopping the crime on the other hand. For example a SAR that made a major contribution to confiscating funds bound for a terrorist organization would score highly on both the seriousness of the crime score and on the level of contribution of the SAR score. By contrast, a SAR that reveals a possible VAT fraud by a local bistro would have a lower score on the seriousness of the crime scale. A SAR that is a false positive might have a zero or even negative score, i.e. useless or worse. The quality metric would be systematically communicated to banks and to a supervisory authority in charge of overseeing the AML system on an end-to-end basis²⁷ in order to ensure that the right balance is struck between AML effectiveness on the one hand, and interference with fundamental rights on the other. The bank would use the feedback to train the TMS algorithms to do better, perhaps using reinforcement learning to train the algorithm to look extra hard for terrorist-related funding, for example. Providing detailed law enforcement feedback to banks may create legal issues, because criminal investigations or national security matters are covered by secrecy. Several solutions may exist, one being to designate an employee in the bank who has national defense secret certification. In France this approach is already used for exchanging sensitive cybersecurity information between banks and cybercrime teams in the government. Another approach would be to provide detailed feedback to certain members of the supervisory authority, and less detailed feedback to the banks.

5.3 Problem 3: lack of "fair balance" because AML processes conduct "general and indiscriminate" analysis of all transaction data.

Highlighted in the CJEU's *Tele2 Sverige – Watson* case, this problem seems difficult to cure for any AML system, whether or not it uses AI. It also poses a Catch 22-style dilemma: the CJEU says that general and indiscriminate processing is prohibited, and that processing of personal data for law enforcement should be targeted based on risks. However in AML the only way to focus on risks is to conduct a risk assessment to define the risky customers areas, which in turn requires processing of all customer data. To overcome this dilemma, it may be possible to argue that the purpose of the systematic processing is in the first instance to identify the risky clusters that merit close monitoring, and that those risky clusters are then subject to more intrusive monitoring in a second step. This approach would be consistent in spirit at least with the risk-based approach in the AML Directives and in the CJEU's *Tele2 Sverige-Watson* case.

5.4 Problem 4: lack of transparency and explainability.

Problem four involves two aspects. The first applies to current AML systems, and involves the prohibition of informing bank customers when they are targeted by a SAR. This problem plagues many surveillance systems, particularly in national security and anti-terrorism surveillance, where

²⁷On the supervisory authority, see cure to problem 5 below.

secrecy is essential²⁸. One solution to this problem would be to require banks to inform customers of the existence of a SAR after a given period has elapsed after transmission to the FIU, for example three months. To maintain secrecy, the FIU would have to inform the bank within that time with regard to particular SARs that are the object of further investigations. Also, as is the case for national security matters, individuals would have the right to apply to the data protection authority or to another authority to verify whether they are linked to a particular SAR.

The explainability problem will require a variety of approaches. As explained recently by France's banking regulator (Dupont et al. [2020]), explainability is necessary for different audiences and different purposes. If a machine learning algorithm creates an alert, explainability will be needed by human reviewers within the bank who review the alert to decide whether to transform the alert into a SAR. FIUs may need explainability to understand why a particular SAR was generated; banking regulatory authorities who want to make sure that the bank's TMS has no gaps, is robust and auditable. From a proportionality standpoint, explainability is needed for two reasons. First, courts and regulatory authorities need to be able to scrutinize the system to make sure that it does not create illegal discrimination. Explanations would need to prove that the algorithm does not target certain groups of the population more than others based on protected attributes such as ethnic or national origin. Among other things, this would require running regular fairness tests such as those done for facial recognition algorithms. Fundamental rights authorities or courts would also have the right to request explanations for individual cases, particularly where an individual files a complaint. Second, individuals need the ability to understand the reasons for a particular decision (such as why the individual was placed in a particular cluster) in order to challenge an individual decision, or the process as a whole in court. This kind of explainability was required both in the Netherlands social security welfare fraud case reviewed above, and in a United States federal court decision involving the scoring of teachers²⁹. In both cases, explainability was a constitutional requirement. Numerous technical solutions are being proposed to provide both "global" and "local" explainability for machine learning models, including explainability by design (Dupont et al. [2020], Beaudouin et al. [2020]). These solutions are likely to bring a cure to this particular fundamental rights problem.

5.5 Problem 5: the absence of an independent supervisory authority with a 360° vision of the system.

As noted above, the current institutional oversight for AML generally does not extend to the law enforcement side of the equation. In equivalent situations involving state-ordered surveillance, an independent commission must have power to oversee the entire chain of processing and report its findings to parliament. In some cases, an approval of the commission would be necessary before deploying certain technologies. To cure this problem for AML, current AML laws could be modified to give specific oversight powers to the data protection authority (for example) to evaluate the whole AML process, going from the design of the TMS by banks to the methods used by law enforcement authorities to follow up on SARs and prosecute money-launderers. The authority's job would be to apply the proportionality test on an ongoing basis, and requiring modifications to the TMS in appropriate cases, or ordering changes in the quality metrics used by law enforcement authorities. If certain AI-based techniques do not prove effective, the authority would order them to be stopped.

²⁸The advocate general of the CJEU recently concluded that France's laws on surveillance are probably illegal for this reason, because they do not provide for individual notification. Advocate general opinion in CJEU Cases C-511-18 and 512/18, 15 January 2020, point 153.

²⁹Houston Federation of Teachers, Local 2415 v. Houston Independent School Dist., 251 F. Supp. 3d 1168 (S.D. Tex. 2017).

6 Conclusion

There is considerable pressure to improve the effectiveness of AML measures, and AI provides a promising avenue. There are numerous barriers to moving to AI-based systems. One important barrier relates to AI's compatibility with the GDPR and its impact on fundamental rights. The biggest fundamental rights question is whether an AI-based system to detect suspicious transactions could satisfy the proportionality test of the CJEU, particularly in light of the *Digital Rights Ireland* and *Tele2 Sverige Watson* decisions. To answer that critical question, we first addressed the question of whether *current* AML processes would satisfy that same proportionality test. We found that no one has addressed that question in detail, at least not since the *Digital Rights Ireland* and *Tele2 Sverige Watson* decisions. To permit our analysis, we first described AML systems as they currently function, and how AI might be introduced. We then applied the proportionality test to AML systems, and identified five major problems. Most of the problems exist regardless of the presence of AI, but AI makes the problems more visible and urgent. A major problem relates to the inability to measure effectiveness of bank AML systems in actually blocking criminal transactions. Without detailed feedback from FIUs, banks and regulators remain in the dark on the actual effectiveness of transaction monitoring measures, making it impossible to apply the proportionality test. A specific problem relating to AI is the lack of explainability of certain machine learning models. Although providing detailed solutions to these problems is outside the scope of this paper, we nevertheless provide some first thoughts on how the five proportionality problems might be cured.

References

- ACAMS - LEXIS NEXIS. Current Industry Perspectives into Anti-Money Laundering Risk Management and Due Diligence. Technical report, Dec. 2015.
- ACPR & TRACFIN. Publication des Lignes directrices conjointes de l'Autorité de contrôle prudentiel et de résolution et de TRACFIN sur les obligations de déclaration et d'information à TRACFIN. Technical report, Nov. 2018. Library Catalog: acpr.banque-france.fr.
- R. Alexy. Constitutional Rights and Proportionality. *Revus. Journal for Constitutional Theory and Philosophy of Law / Revija za ustavno teorijo in filozofijo prava*, (22):51–65, June 2014. ISSN 1581-7652. doi: 10.4000/revus.2783. Number: 22 Publisher: Klub Revus – Center za raziskovanje evropske ustavnosti in demokracije.
- Article 29 Working Party. Opinion on data protection issues related to the prevention of money laundering and terrorist financing,. June 2011.
- V. Beaudouin, I. Bloch, D. Bounie, S. Cléménçon, F. d'Alché Buc, J. Eagan, W. Maxwell, P. Mozharovskyi, and J. Parekh. Flexible and Context-Specific AI Explainability: A Multi-disciplinary Approach, Mar. 2020.
- L. Dupont, O. Fliche, and S. Yang. Governance of Artificial Intelligence in Finance - Discussion Document. Technical report, June 2020.
- EDPS. Assessing the necessity of measures that limit the fundamental right to the protection of personal data: A Toolkit, Jan. 2017.
- Europol. From suspicion to action - Converting financial intelligence into greater operational impact. Technical report, Sept. 2017. Library Catalog: www.europol.europa.eu.

- M. Grazia Porcedda. SURVEILLE Deliverable 2.4 – Paper establishing a classification of technologies on the basis of their intrusiveness into fundamental rights. Technical report, European Commission Seventh Framework Programme, 2011.
- D. Hart. Supreme Court on EU and ECHR proportionality - back to basics, June 2015. Library Catalog: ukhumanrightsblog.com.
- T. Hickman. The substance and structure of proportionality. *Public Law*, pages 694–716, 2008. ISSN 0033-3565.
- IBM. Fighting Financial Crime with AI. Technical report, May 2019.
- W. J. Maxwell. Better regulation applied to the internet. In *Smart(er) Internet Regulation Through Cost-Benefit Analysis : Measuring harms to privacy, freedom of expression, and the internet ecosystem*, i3. Presses des Mines, Paris, Sept. 2017. ISBN 978-2-35671-494-7. Code: Smart(er) Internet Regulation Through Cost-Benefit Analysis : Measuring harms to privacy, freedom of expression, and the internet ecosystem.
- Jun Tang and Jian Yin. Developing an intelligent data discriminating system of anti-money laundering based on SVM. In *2005 International Conference on Machine Learning and Cybernetics*, pages 3453–3457 Vol. 6, Guangzhou, China, 2005. IEEE. ISBN 978-0-7803-9091-1. doi: 10.1109/ICMLC.2005.1527539.
- P. Lascoumes and G. Favarel-garrigues. Les banques sentinelles de l’anti-blanchiment : l’invention d’une spécialité professionnelle dans le secteur financier. Technical report, Dec. 2006.
- N. A. Le Khac and M.-T. Kechadi. Application of Data Mining for Anti-money Laundering Detection: A Case Study. In *2010 IEEE International Conference on Data Mining Workshops*, pages 577–584, Sydney, TBD, Australia, Dec. 2010. IEEE. ISBN 978-1-4244-9244-2. doi: 10.1109/ICDMW.2010.66.
- R. Liu, X.-l. Qian, S. Mao, and S.-z. Zhu. Research on anti-money laundering based on core decision tree algorithm. In *2011 Chinese Control and Decision Conference (CCDC)*, pages 4322–4325, May 2011. doi: 10.1109/CCDC.2011.5968986. ISSN: 1948-9447.
- G. Paravicini. Europe is losing the fight against dirty money. *Politico*, Feb. 2018.
- E. L. Paula, M. Ladeira, R. N. Carvalho, and T. Marzagao. Deep Learning Anomaly Detection as Support Fraud Investigation in Brazilian Exports and Anti-Money Laundering. In *2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 954–960, Anaheim, CA, USA, Dec. 2016. IEEE. ISBN 978-1-5090-6167-9. doi: 10.1109/ICMLA.2016.0172.
- A. Portuese. Principle of Proportionality as Principle of Economic Efficiency. *European Law Journal*, 19(5):612–635, 2013. ISSN 1468-0386. doi: 10.1111/j.1468-0386.2012.00611.x. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0386.2012.00611.x](https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0386.2012.00611.x).
- J.-M. Sauvé. Le principe de proportionnalité, protecteur des libertés, Mar. 2017. Library Catalog: www.conseil-etat.fr.
- D. Savage, Q. Wang, P. Chou, X. Zhang, and X. Yu. Detection of money laundering groups using supervised learning in networks. *arXiv:1608.00708 [physics]*, Aug. 2016. arXiv: 1608.00708.
- TRACFIN. Rapport Annuel d’activité TRACFIN 2018. Technical report, 2018.

- S.-N. Wang and J.-G. Yang. A Money Laundering Risk Evaluation Method Based on Decision Tree. In *2007 International Conference on Machine Learning and Cybernetics*, pages 283–286, Hong Kong, China, Aug. 2007. IEEE. ISBN 978-1-4244-0972-3 978-1-4244-0973-0. doi: 10.1109/ICMLC.2007.4370155.
- 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. *arXiv:1812.00076 [cs]*, Nov. 2018. arXiv: 1812.00076.
- 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. page 7, 2019.