A Layered Approach to Anti-Money Laundering

Author: Lars-Ivar Sellberg, Executive Chairman Scila AB

Introduction

The ever increasing sophistication and creativity when it comes to money laundering can only be matched by the massive volume of such activities. This means that organizations faced with the task of countering money laundering face the complex task of not only having robust AML procedures in place but also being able to maintain their normal core business without being impacted by excessively costly and complex AML solutions.

The somewhat vague term “organization” was used intentionally since there is such a wide range of entities that need to implement AML procedures. Most people would think of banks, but AML is equally important for brokers, trading venues or energy producers to name just a few.

The consequences of failing to prevent money laundering has been made clear in a number of public affairs during recent years resulting in large financial penalties as well as loss of trust.

Long gone are the days in which AML procedures could be met by traditional rules based systems or manual procedures. These methods could actually yield reasonable good results albeit at a very high cost in terms of manual labour, performing checks not automated, weeding out false positives etc. The main problem with these methods is however that they risk becoming stale and ineffective in an environment where the transgressors are highly adaptive and creative in inventing new methods to circumvent existing systems.

The author of this article argues that the solution lies in a multi layered approach, each layer with its own pros and cons but together forming an effective and cost efficient solution. Each layer is discussed individually but also as part of the holistic solution. Rationale and techniques are subjects which are discussed.

It should be noted that while in this article a clear distinction is made between the layers, in a real life software implementation an overlap typically exists between them.

Rule based

A rule based system consists of a number of predefined rules which define suspicious patterns to look for. An example of such a pattern could be multiple deposits just below a certain threshold. The main advantage of an explicit rule like this is that it is well defined what one is looking for.

This makes it easy to explain to a regulator what is looked for. A prerequisite for this of course is that the rules are well documented in a way that can be presented to the regulator. This should include both static as well as dynamic documentation, The static documentation is a description of the rules themselves while the dynamic provides the context when a breach is detected i.e. what were the limits breached, involved parties etc.

Would the regulator provide detailed guidance, although this is typically not the case, such guidance can be directly implemented in a set of rules.

Once a breach is found it is typically quite straightforward, given the proper documentation, to understand why the rule triggered. Note that this does not imply that there are no false positives, on the contrary rule based systems tend to generate quite a few of those and it can be quite time consuming to understand what constitutes a real problem, i.e... it is easy to understand an individual breach, but it can be much harder to see if it is a part of a larger pattern which needs to be investigated.

The main advantage of the rule based approach is also its largest disadvantage, everything that is not explicitly looked for is missed. There is no doubt that this is a significant drawback!

In order for a rule based approach to be a viable solution a prerequisite is that the library of rules is actively maintained and updated as the threat landscape changes. The rules themselves must also be properly parameterized in order to give them at least some flexibility. Backtesting is a key function since adjusted rules need to be applied to existing sets of data in order to find any previously missed suspected transactions.

Using a static set of rules which are not updated is a recipe for disaster and only serves the purpose of giving a false sense of security.

For all its defects the rules set approach is still an important part of most AML software solutions and rightly so. It's simplicity and straightforwardness serve as a first layer of protection. A prerequisite for it to be effective is an active maintenance of the rule library. As a standalone solution it has, as discussed above, some serious drawbacks.

Statistical approach

A statistical approach is where a transaction is compared to a peer group. If deviant then it is flagged as a breach that needs to be investigated. A trivial example might be that the size of a certain transaction is well above the average size as defined for a peer group during the last year.

This approach is less rigid than the rule based implying that one does not run the same risk of missing a suspect transaction because it does not exactly fit a predefined rule. Instead focus is on finding transactions that somehow deviates from its peers, the next step being to investigate whether it actually is a part of a money laundering schema.

A key aspect is how the benchmark is defined with which each transaction is compared. The benchmark is a group of peer transactions that are expected to exhibit the same type of characteristics. They need to be selected from a representative time interval which is long enough to smoothen out any short term fluctuations caused by external factors as changing market conditions. This implies that the length of the benchmark typically is quite long, at 6-12 months.

Besides the time interval, peer transactions are selected from a set that is considered to be similar. Selecting this set can be done in a number of ways depending on circumstances, examples being account types, account owner category, activities, regions to name a few. A key characteristic of any software solution is flexibility when it comes to this

area. Due to the length of the benchmark as well its diversity in terms of composition this can be quite challenging from a technical perspective. It must be possible to change the benchmark, both from a length as well as a composition point of view, ad-hoc in order to try out new scenarios interactively. Having a static benchmark is as bad as a poorly maintained predefined rule library discussed in the previous section.

Once a breach has been detected it needs to be investigated. Here it becomes more complex than the rule based approach since the breach is a result of comparing with a complex benchmark. To assist the investigator a software solution should provide a set of tools, typically visualizations that highlights the breach differences as compared to the peer group. Having a sophisticated set of such investigation tools greatly reduces the time spent to investigate breaches.

As with the rules based approach it is fairly easy to explain to a regulator what is looked for.

A layer using a statistical approach is a vital part of a AML software solution to identify suspicious transactions missed by a purely rule based approach. Used in isolation it might render unnecessary long investigation times however sophisticated the supportive investigative tools are.

Artificial intelligence / Machine learning

There are two main usages for artificial intelligence (AI) / machine learning (ML), breach classification and finding anomalies. For both of these applications the advantage is that there is no need to specifically define what is being looked for. This is in stark contrast to the rule based approach where everything not explicitly looked for is missed. It can not be stressed enough how great an advantage this really is. AI/ML ability to learn from the data set gives it a unique capability to automatically adapt to a changing environment.

This however comes with a cost, as many other good things in life. It is not easily understood what is actually being looked for. This is an inherent problem with all AI/ML algorithms and while there are techniques that attempt to mitigate the problem there is really no good solution. This means that is not possible in a concise and precise manner to explain to a regulator what is actually being looked for. Once a breach is found there are also similar difficulties to understand exactly

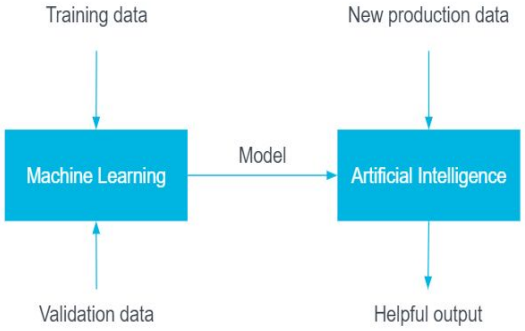

Artificial Intelligence

"Machines or software able to find patterns in data and make decisions based on gained knowledge, autonomously and without predefined rules"

Machine Learning

"Science behind and implementation of how a machine or software finds patterns and gains knowledge from data"

why a breach occurred. Once again there are mitigation procedures but their effectiveness is limited.

This is a drawback which implies that AI/ML as the sole solution for AML software is not really viable at the moment. But what can also be said is that AML software without an AI/ML component is equally non-viable. The ability to automatically adapt and learn from the data set is a powerful ability that really needs to be a part of any serious AML software package.
The purpose of breach classification is to take an already generated breach whether from predefined rule or a statistical analysis and classify it as being an interesting one. In other words false positives are weeded out thereby decreasing manual work otherwise needed to do so.

For breach classification it is advantageous to use supervised learning AL/ML techniques. Supervised learning means that the business user supplies an example set of manually classified breaches which is used as a starting point for the AI/ML algorithm to learn from. New breaches are then automatically classified by the algorithm. As more data becomes available the classification precision is increased.

Note that the supervised learning approach allows the user to define the classification scale. What one user considers to be interesting might differ from another one. Indeed it might be useful to have multiple definitions within one organization.

Non-supervised learning methods are maybe the best option for detecting anomalies i.e. breaches in data sets. As with any method to be effective the supplied data must be rich enough to provide enough context for the AI/ML algorithm.

As with any algorithm AL/ML non-supervise breach detection will yield false positives. An interesting and promising approach is to apply supervised AL/ML classification of breaches generated by non-supervised AI/ML anomaly detection. The combination of the two methods leverages both the knowledge of staff in the supervised part of the method with the generality and adaptiveness of the non-supervised part.

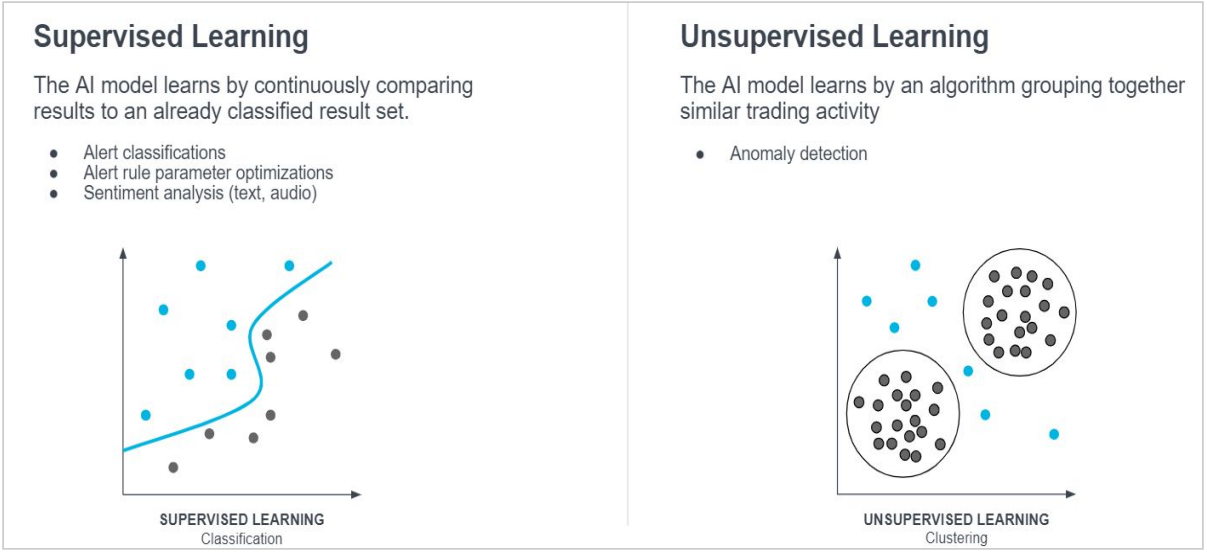

Basing an AML software package on only AI/ML is not a viable solution due to the inherent opaqueness of these methods. At the same time not having an AI/ML part is simply not an alternative considering the unique advantages they offer in terms of flexibility and finding things that would otherwise not even have been looked for.

Summary

There are a number of factors that are important when considering an AML solution:

- Finding relevant breaches
- Minimizing the amount of false breaches
- Transparency, being able to support discussion with the regulator
- Flexibility, the solution needs to stay relevant in a changing landscape
- Cost

While it is possible to achieve reasonably good results with an individual approach, this article argues that the only method covering all the bullets above is a layered approach using several methods simultaneously.

Contact details

www.scila.se
sales@scila.se

SCILA