AI for Third-Party Risk Management

AI for Third-party Risk Management

Going beyond the hype to build a practical approach to machine learning

In many people's minds, artificial intelligence (AI) is associated with self-driving cars or anthropomorphic characterizations like C-3PO from Star Wars or HAL, the murderous computer in "2001: A Space Odyssey." So when vendors start talking about using AI for TPRM, many people envision a scenario in which they give up user control of the system and put the technology in charge of the process. In fact, a practical and strategic approach to AI doesn't replace human intelligence at all; it reflects, refines, and complements it, so that humans are more efficient at making better decisions that result in better outcomes.

This paper is intended to clear up some of the uncertainty about the use of AI in TPRM programs, including what it is, what it can and can't do, and how you can successfully apply it to predict and mitigate third-party risk in your organization.

What is AI?

In general usage, AI simply describes machines, such as computers, that can perform cognitive functions similar to human minds, such as learning or solving problems. To do this, they are able to interpret external data, learn from it, and then apply those learnings to achieve a goal or complete tasks. In TPRM, the goal is to make a decision -- such as whether to onboard a third party – by learning from decisions that have been made in the past, and tasks are the actions that should be taken based on that decision, such as a risk mitigation plan related to a specific control.

Human beings are still responsible for teaching the AI to understand how decisions get made and what the tasks or outcomes should be by showing it examples. The advantage AI has over humans is the ability to efficiently analyze enormous amounts of data, quickly identifying patterns and trends without an explicit decision tree or model. That means that you can show it the input data humans use to make decisions and the output in the form of the decisions that were made. The AI can deduce how to make similar decisions when presented with new input data.

Aside from overuse as a buzzword and Hollywood clichés, one of the reasons AI is a confusing term is because it is actually a broad term for a number of intelligent technologies. Two that are particularly applicable to TPRM are machine learning and natural language processing. Machine learning is the use of computer algorithms that improve their ability to complete a specific task based on experience. When applied to TPRM, that "experience" is the data about all of the decisions your users made in the past and how they continue to make those decisions. The more input (experience) it has, the more accurate and confident the system will be when it comes time to make a decision.

Natural language processing (NLP) is the technology that reads and understands human language. Unlike older technologies such as optical character recognition that recognize individual words/letters or relative position, NLP is able to make connections between words and concepts, such as those used in the free-text responses in risk assessments. It is critical to the way the AI system interprets the data it learns from.

These technologies rely on neural networks, a technology designed to mimic the way human brain cells make connections. The neurons are interconnected layers of algorithms that can be trained to understand the relationship between input and desired output. A simple example would be the evaluation of a risk assessment completed by the third party. A neural network that has been trained using past assessments and their outcomes (approved, denied, conditional, etc.) would be able to learn how risk experts make decisions. When a new assessment arrives, the neural network would be able to apply that knowledge to make or recommend (depending on user configuration) an outcome based on what it's learned.

Applying AI to TPRM

AI is a useful technology in TPRM because it is essentially a decision-making process. Does this third-party align to our risk appetite? Are there any indications that a third party is likely to be involved in bribery or corruption? Are they a critical supplier? AI is really just another way to support that decision making, allowing humans to respond more quickly and with greater confidence and focus on the kinds of activities that AI can't do, like building third-party relationships or conducting site audits.

Those AI-driven decisions can result in three primary kinds of actions:

- Automation
- Intelligence
- Prediction

Automation

Third party risk experts don't need automation to make decisions. In less mature programs they often gather data in spreadsheets and manually review it. However, because of the volume of third-party relationships and associated data, most begin to embrace some stage of automation, which can range from basic to advanced:

STAGE 1

Business Rules

Rules drive the automation process and are based on static thresholds. For instance, if a third party answers "yes" to the question, "Will you be processing PHI as part of your services?", there is a business rule that triggers a data privacy assessment. Rules work well, but trying to create every rule for every contingency in a complex process can be tedious and error-prone.

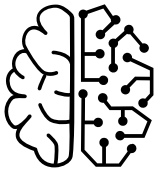

STAGE 2

Analytical Evaluation

A mathematical calculation and evaluation are used to drive automation. For instance, a third party's answers on an assessment are assigned various percentage weights to arrive at a single score. Based on that score, the system determines that the third party is low risk, approves the third party for onboarding, and triggers the onboarding process. Analytical evaluations can also become highly complex to build manually.

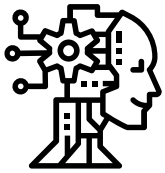

STAGE 3

Machine Learning

The system learns from the actions taken by users and is able to make similar decisions without the need to have an explicitly defined model. For instance, when a third party completes an assessment, the system knows how humans would respond based on an analysis of all of the similar decisions that it has been exposed to.

Configurators can decide (based on individual engine training quality, the criticality of particular decisions, and engine confidence) to allow certain statuses and classifications to be set or changed without risk expert input or review, automating rote or relatively low-risk decision making tasks.

Intelligence

There are some decisions that are not as suited for automation. Perhaps an appropriate decision relies on a combination of data analysis and some kind of human judgement. Or the associated risk is considered too high for an organization to feel comfortable automating it entirely. AI can provide the intelligence needed to augment human decision-making.

In this use case, AI uses accumulated data to advise a risk expert on an attribute, such as classification and/or status selection. When the risk expert is presented with a task to complete, they could also be presented with a recommendation generated using machine learning along with a confidence level. For instance, a third party may have completed an anti-bribery/corruption assessment and been screened against the US Consolidated Screening List. When reviewing the assessment and search results, the risk expert may be given an AI-generated yes/no recommendation along with a level of decision confidence from 0-100% confidence and an evaluation of the model quality from 0-100%. The risk expert still makes the ultimate decision, and the machine learning engine learns from the decision so that it continues to improve the model and the confidence.

An AI that has been taught to recognize good decisions also recognizes bad decisions because they don't conform to the patterns of previous actions. When utilized in this way, an AI-enabled system can alert users to non-conforming decisions, preventing possible risk exposure and regulatory non-compliance. For instance, if a user were to approve a third party that would normally be considered high risk without a mitigation plan, the system could flag the decision for review by a human.

Prediction

Many TPRM programs use continuous monitoring to detect changes that would indicate when the risk profile of a third party is likely to change. Sometimes these changes are large and obvious, but often early signs may be more subtle and based on more than one key risk indicator (KRI), making it easy for busy TPRM teams to miss early warning signals. In other cases, monitoring can generate too much data, overwhelming those responsible for filtering through numerous alerts with varying degrees of relevance.

Using AI, when input data changes sufficiently that a particular classification or status should change based on machine learning intelligence, the system notifies risk experts to review the change and may automatically trigger a remediation process.

For example, a TPRM system may use continuous monitoring to regularly check individuals and corporations against the US Consolidated Screening List for ABAC compliance. If a screening were to come back with an unexpected change in status that indicated a higher risk, an AI-enabled system could be configured to change the status from "Compliant" to "Non-compliant," just as a risk expert would have done in the same circumstances. This change could also trigger an Issue & Corrective Action workflow to alert humans to the issue.

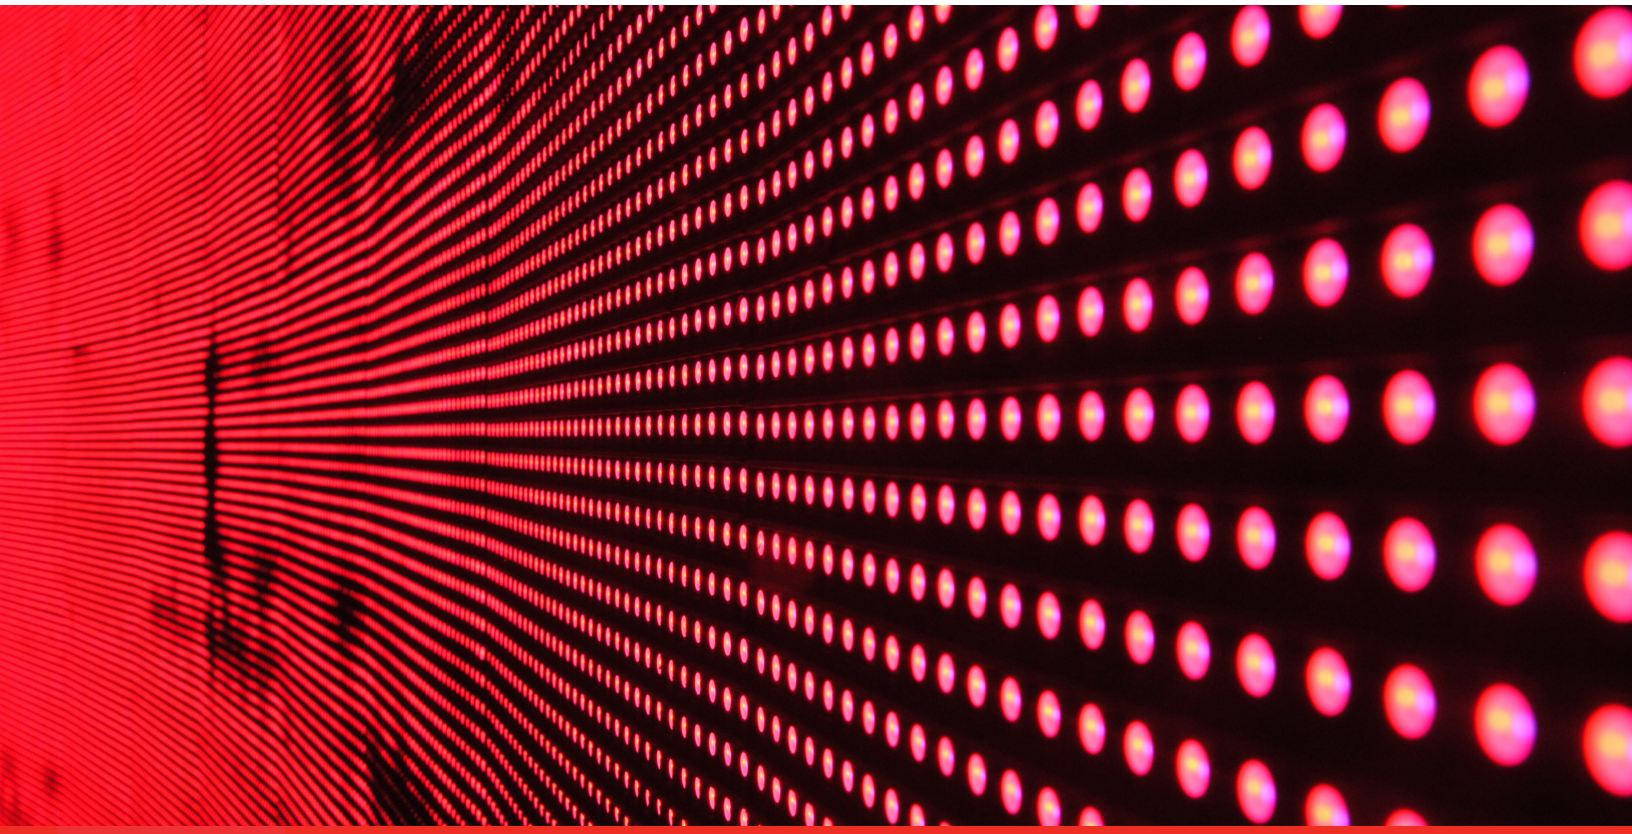

Benefits of AI for TPRM

Machine learning and natural language processing address many of the challenges that TPRM programs face, such as:

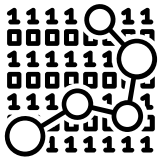

DEALING WITH VAST AMOUNTS OF DATA

Even with highly automated systems, TPRM teams can often be overwhelmed with the amount of data they have to manage, from assessments and related documents, to continuous monitoring data, to external risk intelligence data. The sheer volume of data increases the risk of human error, such as missing important details or subtle changes that can be detected by a machine learning engine.

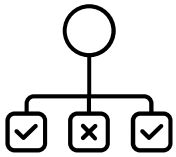

INCREASING CONFIDENCE IN DECISION-MAKING

By supporting and validating decision-making processes, machine learning instills more faith in the decision-making expertise of the first line of defense (the risk owner) and the second line of defense (the risk expert). That confidence, in turn, improves the activities of the third line of defense (audit).

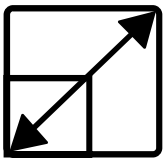

SCALABILITY

A recent survey found that more than 40% of TPRM teams have only 1-5 people, yet roughly half of these small teams manage more than more than 500 third parties, including more than one-third who were managing more than 5,000. Growing reliance on third parties, increasing and evolving regulations, and increased C-level/board focus on risks like cybersecurity and supply chain resilience increase the pressure on TPRM programs despite the fact that that the majority of programs don't feel they are adequately resourced. AI can help teams work more efficiently and scale their programs to meet these demands with greater quality and fewer mistakes.

REINFORCING ADHERENCE TO ORGANIZATIONAL BEST PRACTICE

Because a practical AI approach will use your organization's decisions to train the system, your organizational expertise is trained into the system. There are systems that attempt to use aggregated data from a buyer community to train machine learning, but TPRM is not a use case that is suited to a one-size-fits-all approach. Most organizations use their own assessments, have their own risk appetites, and make decisions on specific priorities, requiring that machine learning be trained on their individual data.

What AI Isn't

The hype around AI and machine learning often gives the impression that the technology can solve any problem, but it's not infallible. Without a framework and strategy driven by human intelligence, AI/machine learning projects will fail. Often these failures are driven by common misconceptions, so it's important to understand that this technology is NOT:

- **A replacement for defining a process**. Some people fear that AI will affect their control of the system; others hope that AI will relieve them of the need to build a governance framework. Ultimately, the organization still has to define the process. It doesn't have human intelligence or experience. Machine learning understands patterns and trends and how to relate them to actions, but it has no innate understandings of concepts. Humans have to show the AI system reasonable decisions first in order to have a process that aligns to program requirements. Without putting in the strategic human intelligence, you will fail.

- **A solution for bad past decisions**. Not only do you have to know the decisions you need to make, you have to be able to train the engine on how good decisions are made. That means that the past decisions you use to train the system have to make sense. AI won't be able to miraculously fix bad decisions, so it's important to maintain a data set that demonstrates your defined process as you'd like it to operate.

- **A workaround for bad data**. One of the dangers of implementing a machine learning solution is rushing in before addressing data issues, such as data quality. Normalizing the data is critical to the success of any machine learning project in order for it to find patterns and make decisions that make sense, so it's important not to undertake an AI project until any data issues have been addressed. However, if the bulk of your data is fairly reliable, an AI-enabled system can help to identify outliers that may represent training opportunities or other weaknesses in the process.

- **Superior to humans**. It cannot be overstated that there is a danger in putting blind faith in AI systems. They don't have the inherent understanding that humans do. Organizations implementing AI still need the natural intelligence of smart people who know and support your business, which cannot be replaced by any vendor's machine learning solution.

TPRM program leaders should be evaluating AI support for decision making, as it will become increasingly important. Without it, many programs will find it difficult to efficiently synthesize the overwhelming amounts of data from inside and outside of the system and continue to focus on the tasks that are a priority for human intervention. And while it won't be as personable as a Hollywood depiction, a realistic approach and the work you do to make sure it's successful determines whether you end up with the equivalent of a helpful droid or an obstructive menace.

The Aravo Decision Engine

The Aravo Decision Engine is a configurable natural language processing machine learning platform that is built into the larger Aravo business process automation platform. As an inherent capability of the Aravo Platform, the Decision Engine is available across all Aravo implementations, including ready-to-use applications for:

- Third Party Management
- Information Security
- Data Privacy
- Anti-bribery/Corruption
- Financial Services Risk Assessment

Aravo clients can leverage the Decision Engine for any use case in which users of the platform are making decisions based on review of particular sets of input data. These decisions made over time are used to train the Decision Engine to advise on or completely automate that decision-making process.

Why Aravo?

In The Forrester Wave for Supplier Risk and Performance Management, Q3 2020, Aravo was recognized as "an SRPM leader thanks to its domain expertise and AI vision." The report noted that "Aravo is ahead of its competitors in applying AI to streamline risk assessment and monitoring."

Aravo customers benefit from a unique combination of 20 years of experience in delivering solutions to the world's largest brands and developing award-winning technology to gain:

VALUE

Because you can configure the Aravo Decision Engine to reflect your requirements, data, and decision-making by leveraging your users' past decisions, you can quickly drive the outcomes that are important to your organization.

EASE OF IMPLEMENTATION

As native functionality that is an inherent component of the Aravo Platform, there is no prolonged deployment. When the functionality is activated in the system, configuration users can immediately use a simple interface to begin training the Decision Engine on specific tasks.

SUCCESSFUL OUTCOMES

Even though the Decision Engine is configurable, you have the complete support of an experienced team of TPRM experts that are responsible for Aravo's unparalleled track record of customer success and committed to help you achieve your goals.

Learn more about Aravo and request a demo at www.aravo.com

[1] Taking the Pulse of Third Party Risk Management 2020

The Definition of Better Business

Better business is built on acting with integrity.

It commands better performance, delivering better efficiency, collaboration, and financial outcomes.

It inspires trust.

But better business is more than that – it's about lifting the ethical standard of an entire business ecosystem to build a better world.

For More Information


Visit us at aravo.com


Email us at info@aravo.com

Call us at:


+1.415.835.7600 [US]

+44 (0) 203-743-3099 [EMEA]