WHITE PAPER

# What is next-generation AML?

The fight against financial crime fortified with robotics, semantic analysis and artificial intelligence

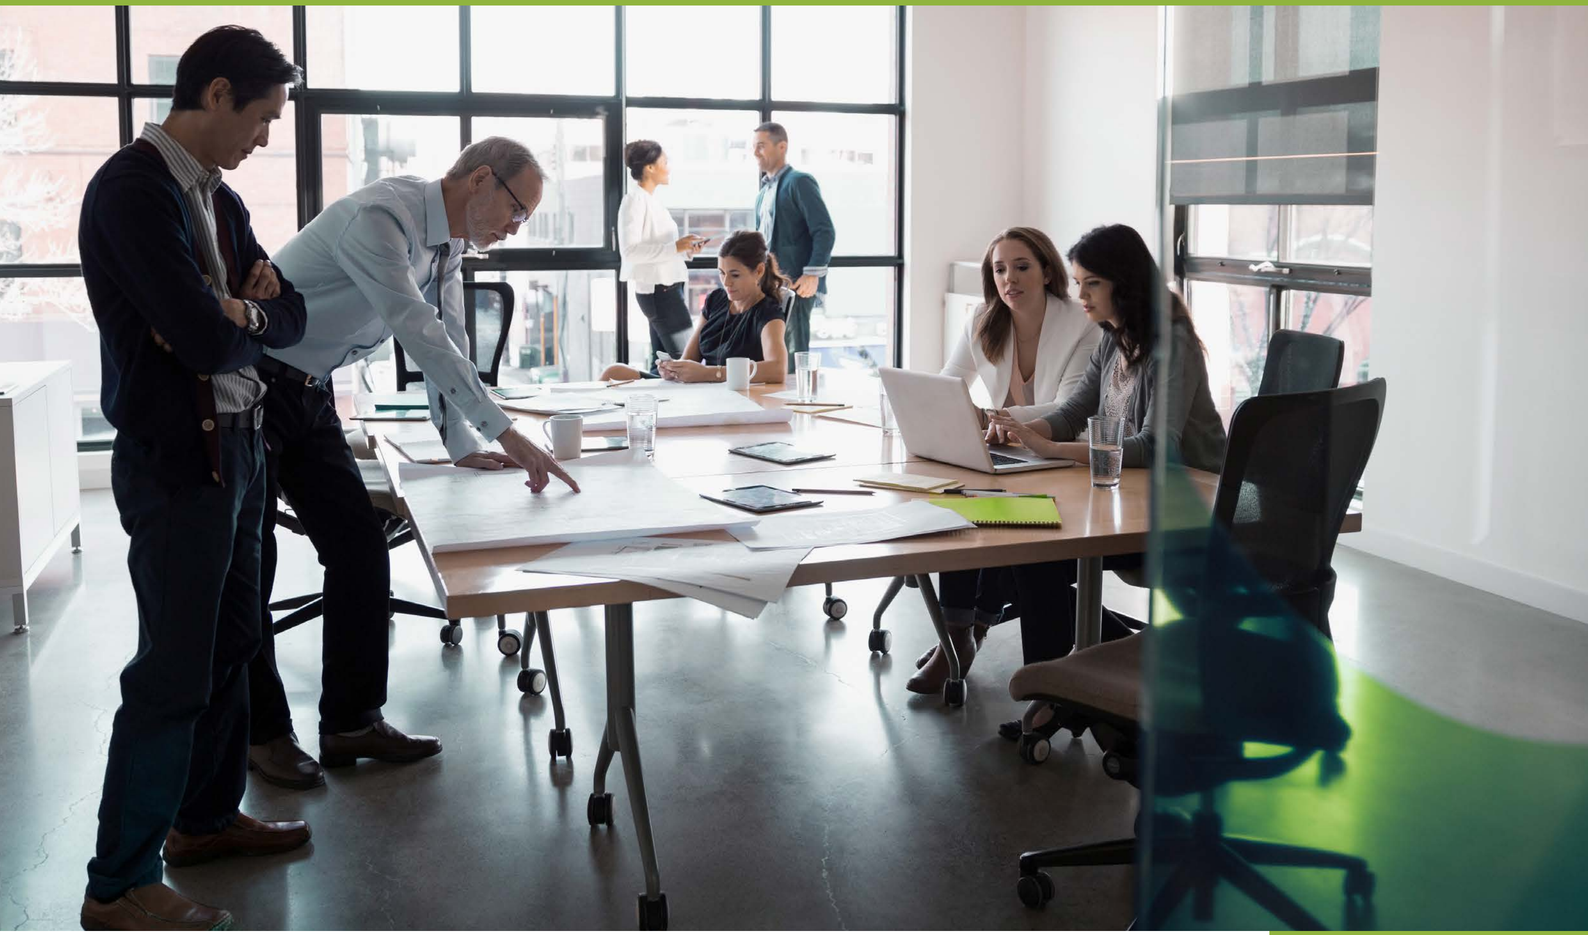

# Contents

Featuring:

**Wallace Chow,** AML Practice Lead, FCC Analytics

**Amith Satheesh,** Principal Solutions Architect and Global Lead, AML Analytics, Security Intelligence Practice, SAS

**Beth Herron**, Senior Solutions Architect and Americas AML Lead, Security Intelligence Practice, SAS

**David Stewart**, Director, Security Intelligence Practice for Financial Services, SAS

# The drive to advance the state of AML

Drug trafficking, smuggling, fraud, extortion and corruption – all illegal but also enticingly lucrative. Proceeds from these criminal activities represent an estimated 2% to 5% of global GDP. That's equivalent to US$800 billion to $2 trillion a year, according to the United Nations Office on Drugs and Crime.[1]

Money laundering disguises the sources and destinations of these funds and fuels some dire downstream effects, such as compromised financial systems and the means to keep terrorists and crime rings in business.

Anti-money laundering (AML) has been a hot topic – and an intensifying regulatory pain point – for financial institutions for decades. For example:

- The USA PATRIOT Act, swiftly enacted after 9/11, expanded requirements for detecting and reporting activities that could signal money laundering or terrorist financing.
- The New York State Department of Financial Services regulation 504, "The Final Rule," effective in 2017, added more granular and stringent control expectations for anyone operating under New York Banking Law, from multinational banks to local check cashing outfits.
- The Fifth EU Anti-Money Laundering Directive (5AMLD), which takes effect in January 2020, expands the scope of covered entities, introduces stricter due diligence and disclosure requirements, and puts the onus on European enterprises to align with US regulations.

Financial institutions have addressed these ever-expanding AML business requirements with four basic types of software:

- **Transaction monitoring systems** that flag transaction patterns that could indicate suspicious activities.
- **Currency transaction reporting systems** to report large cash transactions ($10,000 and more in the US).
- **Customer due diligence/know your customer systems** to get clarity into customer relationships and risks.
- **Watchlist screening** to identify suspicious or sanctioned individuals and organizations.

Now there's a lot of talk about advancing the anti-money laundering arsenal to the next level, sometimes referred to as next-generation AML, AML 2.0 or AML 3.0. Whatever you call the next wave of AML technology, it's about solutions that draw on such advances as robotics, semantic analysis and artificial intelligence (AI).

It's about making AML processes more efficient and effective. And it's about augmenting the traditional rules-based approaches to drive down the rate of false positives and more accurately detect activity worth investigating.

[1] United Nations Office of Drugs and Crime, https://www.unodc.org/unodc/en/money-laundering/globalization.html, accessed March 10, 2019.

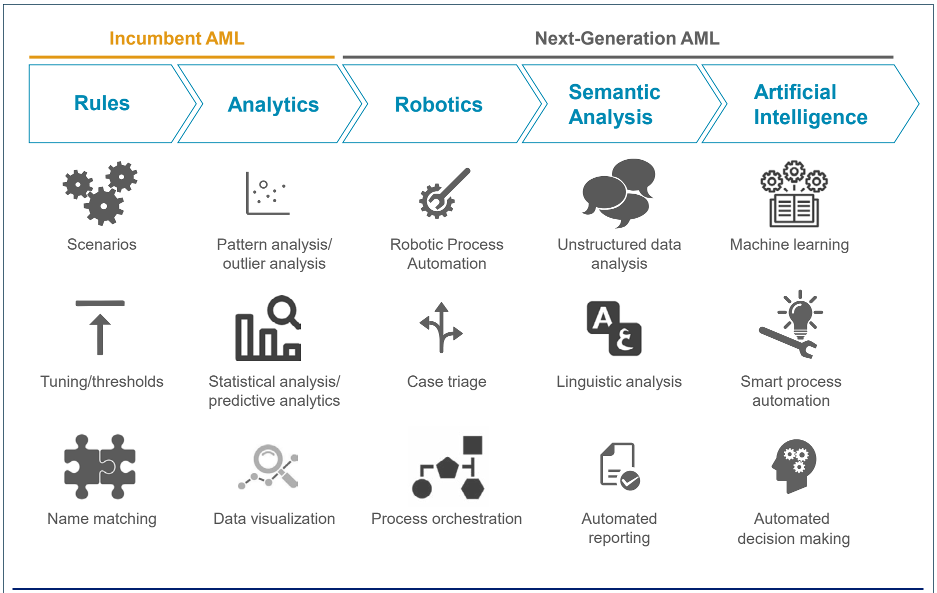

Source: Celent

Figure 1: The evolution of AML compliance technology.

As with everything in IT, the evolution to next-generation AML is coming in phases, starting with easy wins and advancing to proof-of-concept testing for more advanced capabilities.

"Much of the work in the last 18 months has been to apply artificial intelligence to some low-hanging fruit, such as using robotic process automation to investigate and prepare cases more quickly," says David Stewart, Director of the Security Intelligence Practice for Financial Services at SAS. "As of 2018, we're starting to see adoption of machine learning not just for process automation, scoring and hibernation, but to supplement or even replace traditional Boolean logic for detection of potentially suspicious activity."

# The state of next-gen AML around the globe

The pace of this AML evolution varies across global regions – faster where regulatory oversight is more stringent, slower elsewhere. No surprises there. But given the ease, speed and complexity with which digital currency flows across international borders, there's serious interest everywhere in AI-powered AML approaches.

## Asia Pacific: In discovery, with growing interest but few guidelines

"There is very dramatic growth in AML not just in the financial center of Hong Kong but in the whole Asia Pacific region," says Wallace Chow, AML Practice Lead at FCC Analytics. "Many SAS customers are eager to adopt next-generation AML solutions with artificial intelligence both to automate processes and improve detection." Government entities are asking banks to adopt more analytical approaches to scenario tuning, threshold setting and other components of overall AML programs – but they're not offering up guidelines for doing so, yet.

## Europe: Gaining ground but still grappling with basics

The AML evolution is further along in Europe, but not where it could be, says Amith Satheesh, AML Analytics Lead at SAS. "There's obviously a lot of emphasis on monitoring and detecting suspicious activity, especially around the Baltic states. Artificial intelligence and machine learning have been the buzzwords lately, influenced by European banks that have some presence in New York or other parts of the US. Many are moving from an early discovery stage to more of an application and implementation stage, and real use cases are out there producing results."

But overall, the efforts in Europe have been mostly reactive. Much of the work still revolves around the challenges of getting the basics right, and that puts Europe about 10 years behind the US, Satheesh says. "Many banks are in catch-up mode. They're still focusing on foundational areas such as segmentation, single view of the customer, moving from account-based monitoring to customer-based monitoring, threshold setting, building better-performing scenarios, reducing false positives and addressing data quality.

"That's where more of the focus is right now, rather than on implementing predictive models to build better detection scenarios. Progress is being made. There are discussions and active proof-of-concept efforts to achieve those results, and people are open to hearing about new AI and machine learning modes."

## US: Driven to maturity sooner in a stricter regulatory environment

In the US market, there has been an exponential growth in AI and machine learning in AML over the last 10 years, especially in areas of transaction monitoring, says Beth Herron, who leads the Americas AML team at SAS. The "final rule" era has put the spotlight on having model risk rigor around AML programs, which has created a culture of compliance and pushed the US to greater analytic maturity than other markets.

"Folks are paying a lot of attention to detection mechanisms and questioning whether their rules and scenarios are really capturing all the risk at their institutions," says Herron. "So we're seeing a lot of experimentation in sandboxes or analytic environments, a lot of proof-of-concept projects, and it's encouraging that we're seeing a lot of those projects moving past the experimentation stage and into production, with great results."

> "Some of the more advanced early adopters of AI are getting those pilot projects over the line, and doing so with great results."
>
> **Beth Herron**
> Americas AML Lead
> Security Intelligence Practice, SAS

# Ten keys to success with next-generation AML

## 1. Innovate, but with caution

In December 2018, five federal financial authorities in the US, including the Board of Governors of the Federal Reserve System, issued a joint statement calling for "innovative efforts to combat money laundering and terrorist financing." The three-page statement contained some mixed messages and light-handed assurances. Right up front, it encouraged financial institutions to "consider, evaluate, and where appropriate, responsibly implement innovative approaches" to meet AML compliance obligations.

"So, while we have this tone of encouraging innovation to drive our processes forward," said Herron, "the word 'responsibly' sends a clear message that they expect this to take place in a sandbox or run in parallel for some time, to truly make sure we're getting the results that we would expect, and these new techniques are stable and explainable."

The statement says these exploratory efforts "should not subject banks to supervisory criticism even if the pilot programs ultimately prove unsuccessful." Banks would probably prefer more definitive language than 'should not,' especially since AI-powered methods will probably find more suspicious activity and potentially make existing systems appear deficient.

## 2. Establish rigorous model governance

A comprehensive AML program establishes three lines of defense, says Satheesh. At the first line, you're quantifying the risk and figuring out what you have to monitor. At the second, you're developing those controls. And in the third line, you have to be in a position to challenge those controls. Are the algorithms still working? Are cases appropriately expedited or hibernated? Are models being monitored and tuned as necessary?

"I don't see a big emphasis on that third line," says Satheesh. "Europe could benefit from more rigor in this area. It's highly reactive at the moment. Only when a bank is in trouble and in the news are questions raised about model risk governance."

## 3. Securely share data across borders

In the UK, the Financial Conduct Authority has been holding hackathons in collaboration with advisory firms, trying to devise solid ways of using AI and machine learning in AML use cases, addressing the complexities of international transactions.

"What has come out of some of those sessions is consensus that before we can adopt machine learning in a meaningful way, we need to have a way to securely share data across borders in a GDPR world," says Stewart. "We need technologies like homomorphic encryption, which would allow us to transfer information across borders." (Homomorphic encryption – still largely a theoretical capability – is a form of encryption that allows computations to be performed on data without decrypting it.)

## 4. Consider a hybrid approach

As regulators have made clear, traditional ways of monitoring transactions against rules are not enough, for several reasons. Rules take too narrow a view, a hindsight view based on what we know about past patterns. Rules are relatively easy to circumvent yet unwieldy to maintain.

But nobody is ready to abandon their rules-based systems and fully replace them with analytical models and robotics. "We're seeing a hybrid approach," says Herron. Use rules where they do the job, models where rules would fall short. "There are some areas where patterns are so well defined that there's no need for a model, or there's the expectation that this is a covered area, and maybe the target areas or the outcomes are so rare that it's difficult to fully automate."

Analytical models shine in situations that call for discerning complex patterns from well-defined behaviors. For instance, alert scoring and hibernation models can clarify the risk associated with a package of alerts, rather than just looking at a single transaction.

"This is a way to get more value out of a traditional transaction monitoring system," says Herron. "Maybe one wire transaction doesn't necessarily look risky, but when you see it combined with additional activity, that is going to raise suspicion. So we're seeing a move away from looking at individual activities and more toward a behavioral profile or multiple activities that can really drive a successful case."

## 5. Take a hard look at the data foundation

The underlying truism of any form of data science is "garbage in, garbage out." That mantra takes on new meaning with machine learning techniques, because the system learns from the data and can therefore auto-generate even more garbage as it goes. "There's no point performing AI and machine learning on bad data; you're not going to get anything out of it," said Satheesh.

Many financial institutions haven't evolved much from a data perspective in the last 10 years, says Satheesh, at least from his vantage point in Europe. "Many banks still do not have a single view of the customer, for instance. Scenarios are continuing to generate alerts at the account level, and all of them are being worked by different analysts, without a full view of the customer from a risk perspective."

But big banks have it mastered, right? "Larger banks, especially if they've grown through mergers and acquisitions, have more silo systems and databases feeding into a compliance data hub, and not all of them are connected," says Satheesh. "I see a lot of data quality issues, not being able to draw the line between those source systems all the way to the destination. It's a big challenge."

## 6. Take a granular approach to thresholding

Get more rigorous than simply dividing personal and commercial clients. That's a start, but once you've done that basic segmentation, you want to get more granular to reflect your business model. You could take a basic approach, such as taking standard deviations and bucketing organizations into small, medium and large based on their total transactions. Or you could create segments by product or transaction type.

Or you could take a more sophisticated approach. K-means clustering – a popular unsupervised machine learning algorithm – makes inferences from the data based on input variables, without referring to known outcomes. This analytical technique helps you understand how variables interact and naturally clusters entities into different groups. Once you have accomplished that, you can zoom in on your scenarios and look very closely at risks specific to those clusters.

Getting more granular works. For example, for one SAS customer using this approach, segmentation refined by machine learning improved productivity rates from 2.8% to 6.8% while decreasing false positives and overall alert volume. After tuning the thresholds for each segment, the productivity rate – the percentage of alerts that resulted in cases worth investigating – jumped to 10.4%.

## 7. Focus on what matters, defer what might matter later

Why expect investigators to look into every alert that pops up? Some will represent only scant risk. Couldn't you safely put it on hold until it rates a higher risk score? Yes. Auto-referral or hibernation functions tap into a broad range of relevant information to either expedite or hold off on escalating an alert for review.

"Once the alert has been generated, an AI engine helps calculate the risk score to see if the alert is actually good for investigation," said Chow. Those with a very low score – as determined from multiple risk variables and categories – could be put in hibernation until the score warrants a closer look. "This saves quite a lot of human effort and ensures that investigators get good alerts to be investigated, those with a high likelihood of converting to SARs."

Don't chase what may turn out to be false positives. Low-risk alerts can be put in hibernation mode while the system gathers more information related to the case and triggers an alert when additional relevant information is found.

## 8. Use machine learning to detect rare events

Through unsupervised learning, a model can analyze a large amount of data and identify hidden patterns and find things that are out of the norm. "In this case, you didn't necessarily know if people are good or bad; you're just looking for those who represent 'edge cases,' people who are behaving out of norm, relative to their peers," says Stewart.

In one case, SAS® used supervised learning to find a handful of illicit players hidden in 1.7 billion transactions, says Herron. "It would be incredibly difficult to apply rules and knowledge to this task. In this context, we really needed to use some advanced methods.

"As our target variable, we used the behaviors of a specific population of customers we knew to be registered with FinCEN. The model then looked for those behaviors in the rest of the population to spot those who were not registered as the type of entity we were seeking but behaving like they are." Rather than a handful of illicit players, the model found dozens.

"When we think about forensic teams looking for that proverbial needle in the haystack, these techniques can be very, very powerful for churning through large volumes of data to find that piece of risk that's difficult to find using traditional methods," says Herron.

## 9. Embed best practices into reusable packages

“As we learn from pilot projects, we have packaged up a lot of those best practices into more of a packaged bot we’re calling an adaptive learning and intelligent agent system,” said Stewart.

This capability automates the creation, publishing and retraining of machine learning models. It offers data scientists and compliance analysts pre-selected variables and recommends best-fit models based on sampling of rare events. In short, it enables more meaningful data analysis with less effort from data scientists – and helps speed the detection of suspicious activities.

## 10. Converge or at least integrate financial crimes systems and processes

Financial institutions today have various risk functions – fraud, public security, cyber-security, credit risk, AML – typically with disparate systems, people and processes.

“I foresee the convergence or integration of all these different financial crimes functions, with a more comprehensive and fully resolved workflow across them,” says Satheesh. “Instead of looking at a customer in an AML lens or a fraud lens, the future will be looking at customers holistically through all these lenses in one central place.”

Data orchestration, analytics development, decisions, case management, reporting and governance will take place in a unified environment that supports bidirectional communication and synergies among functions.

“When that happens, we’re going to see all of these risk functions be more powerful, more efficient and more cost-effective,” says Satheesh.

> “Right now we are experimenting, doing proof-of-concepts and on-the-side implementations. But in the future state I do see this as more of a plug-and-play kind of feature with a solid analytics foundation.”
>
> **Amith Satheesh**,
> AML Analytics Lead
> Security Intelligence Practice
> SAS

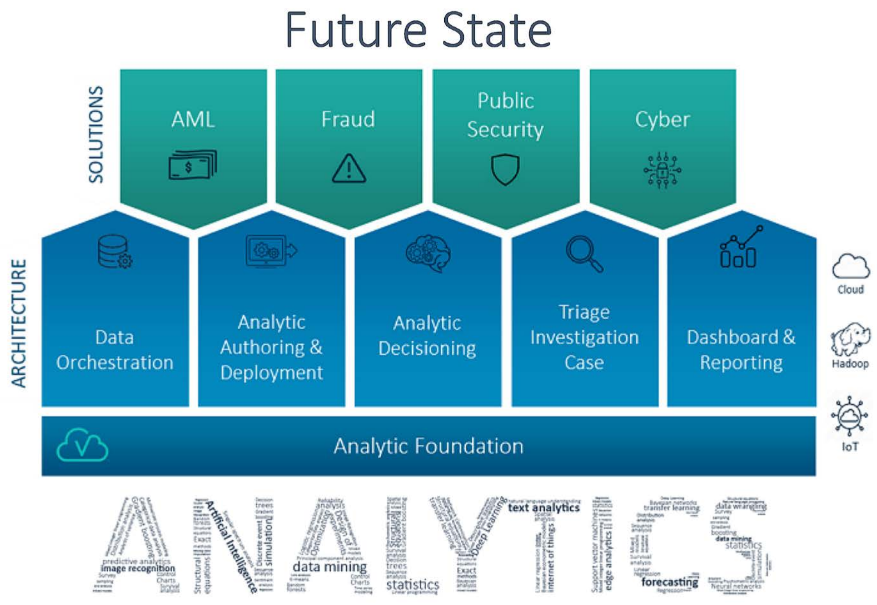

Figure 2. The future state of financial crimes and cybersecurity.

# Advanced AML in action: Use cases around the world

## Intelligent automation

A retail and commercial services bank in the Asia Pacific region faced several challenges. Transaction volume had grown so rapidly that it was impossible to manually review all alerts. A high false-positive rate further stretched investigators' capacity. The bank wanted to automate the review of low- to moderate-value work items and use analytics to detect hidden risks.

SAS developed an alert scoring and hibernation approach for predicting which cases would lead to productive investigations. The scoring algorithm was based on transaction risk, entity risk, network risk, scenario risk and customer risk. Applying ensemble methods using gradient boosting and deep neural networks, the model fully automated the alert review process and reduced false positives by 33% – for significant savings in costs and investigators' time.

## Rare event detection

A Tier 1 global bank wanted to test the validity of applying AI to detect hidden risks and false positives within its client base. The bank knew its rules-based transaction monitoring system was biased toward known risks, and that regulators expected to see the bank expand its coverage over high-risk businesses.

The bank engaged SAS to mine a large data set and find something they didn't know. Specifically, they wanted to ask, were there high-risk customers in their customer base that weren't properly classified?

"We applied some of our newer algorithms – a random forest model with 200 trees," said Stewart. "The model split out what we deemed to be customers who were behaving as money service businesses but who were not declared as such during the onboarding process, nor were they legally registered as money services businesses.

"In a population of nearly 2 billion transactions (aggregated in about 10 minutes), we found 416 suspected money service businesses, 89 previously unknown or unregistered MSBs, which, through further triage, resulted in dozens of productive cases."

## AI-powered detection

A Tier 2 regional US bank wanted to modernize its incumbent rules-based AML transaction monitoring system. Its library of 200 transaction monitoring scenarios was difficult to manage. The bank needed to reduce the volume of low-value events, address gaps in coverage and improve SAR conversion rates.

The bank deployed a SAS neural network model into production to replace 10 cash activity scenarios from their transaction monitoring system. The data scientists noted that the model uses 75 to 80 variables to detect potentially suspicious cash activity, whereas their previous rules-based scenarios used only six to 12 variables.

"They have tripled their SAR conversion rates and cut their monthly work items by 50%," said Stewart. "Their plan is to roll off 200 other current rules-based scenarios into maybe 25 or 40 analytics-based scenarios for certain types of high-risk typologies, and over the next 18 months move everything else to machine-learning strategies."

### Natural language processing

A Tier 1 global bank wanted to improve the customer experience for trade finance business by accelerating due diligence reviews and improving accuracy. As it was, staff couldn't keep pace with service level agreements for trade document review. It was time-consuming to ensure packages passed compliance review. Automation would reduce the burden on employees and reduce costs.

The bank deployed deep learning algorithms to classify the types of documents under review. It applied contextual analysis to specific classes of documents, such as identifying under- or over-invoicing. In this case, automation driven by machine learning reduced the effort from two weeks of staff time to less than a minute.

## Closing thoughts

Next-generation AML is coming to the forefront as the financial services industry goes through massive digital transformation and as regulators keep upping their definition of "reasonable" control and governance. Robotics, semantic analysis and artificial intelligence – particularly machine learning – will be central to this evolution.

Instead of simply reacting to past information, machine learning delivers a forward-looking advantage. You can distill new data elements not previously known or available for AML models. Let the computer uncover patterns the human eye would never see. Validate those insights and feed the results back into modeling efforts. The more training a model gets with feedback data, the smarter it becomes and the less tuning it requires.

Not ready to give up your trusted rules? Adopting machine learning doesn't mean you have to abandon established ways of working. The right route for an organization's AML modeling might be a complementary approach, a hybrid of rules and predictive models.

The use cases described in this paper are things SAS has been doing for years with very progressive clients. What has changed is that the barrier of entry has been reduced for smaller institutions. SAS is packaging proven machine learning techniques for AML tasks to automate repetitive, manual processes, more accurately detect potentially suspicious activity, and put these capabilities in the hands of more financial services organizations.

Learn more at sas.com/aml-ctf.

If you're making decisions on only on the structured data fields in your systems, then you're making decisions with less than 20 percent of the available data. To reach the needed level of accuracy, it's critical to bring in unstructured data from sources such as previous SARs, comments in customer service records and recordings of phone calls.

By applying advanced analytics and powerful machine learning on a unified platform, financial services firms can gain a holistic view of risk, uncover more financial crime patterns, reduce false positives and run more efficient investigations.

To contact your local SAS office, please visit: sas.com/offices