Brought to you by the publishers of COMPLIANCE WEEK

INSIDE THIS PUBLICATION:

Artificial Intelligence, Welcome to the World of Compliance

Message from the Editor in Chief

The rate of computing power is growing so quickly that by 2020, we will finally have a computer with the same raw processing power of a human brain. By 2050, we will have a computer with the same raw processing power of all human brains combined. We are seeing the computing power of machines hit a hockey stick moment and, as that power increases, so does the advent of truly artificial intelligence—machines that think on their own and learn on their own.

AI has long been the realm of science fiction, but in compliance, it takes on a new meaning, as AI-driven computer systems can provide organizations with the ability to process huge amounts of data either in search for particular documents or to screen communications for the usage of certain words or phrases, or scanning financial transactions for signs of money laundering or other impropriety. With this kind of increase in capabilities, compliance officers are understandably interested in what AI can do for them. But first, we need to understand better what exactly AI is and does in a compliance context and where we are between the promise of the technology and what it actually delivers.

This e-book gathers a few of Compliance Week's most recent articles on the subject of AI, but this is hardly the end of the discussion. If you like what you see here, be sure to subscribe to Compliance Week and visit www.complianceweek.com to read the latest news, analysis, and commentary on AI and the other technological innovations that are driving the discipline, profession, and industry of compliance into the future.

Bill Coffin
Editor in Chief
Compliance Week

About us

COMPLIANCE WEEK

Compliance Week, published by Wilmington plc, is an information service on corporate governance, risk, and compliance that features a weekly electronic newsletter, a monthly print magazine, proprietary databases, industry-leading events, and a variety of interactive features and forums.

Founded in 2002, Compliance Week has become the go to resource for public company risk, compliance, and audit executives; Compliance Week now reaches more than 60,000 financial, legal, audit, risk, and compliance executives.

Inside this e-Book

Rise of the machines

Artificial intelligence could revolutionize compliance

It may sound futuristic, but "thinking machines" are poised to revolutionize compliance. Artificial intelligence, proponents say, can take care of grunt work, freeing audit and compliance professionals to focus on matters that befit their skills. Advanced automation, however, says **Joe Mont**, isn't without concerns and pitfalls.

The concept of "thinking" machines and artificial intelligence is a familiar science fiction trope, one populated by nihilistic computers and robots with a Pinocchio complex. Computers may not be so fantastically self-aware yet, but artificial intelligence is nevertheless poised to revolutionize compliance.

Already, firms are offering software platforms that promise to automate otherwise routine tasks and improve upon fraud detection audits, anti-money laundering protocols, and know-your-customer screening. The pitch is as simple as the technology is complex: let machines scan through a company's data to do the grunt work of simple investigations, better utilizing the skills and expertise of human personnel.

Advanced automation isn't without concerns and pitfalls. Do cutting-edge technologies fit with legacy systems and existing automation? What of that utmost "legacy" concern: employees? How can, for example, compliance officers maximize the benefits of AI, and have it compliment their efforts, without losing a necessary "human touch"?

A recent study by audit, tax, and advisory firm KPMG concluded that "the convergence of robotic process automation, machine learning, cognitive computing, artificial intelligence, and advanced analytics are driving unparalleled business model transformation."

Less promising, 81 percent of CEOs it surveyed as part of its 2016 U.S. CEO Outlook said they are concerned about having to consider the integration of basic automated business processes with artificial intelligence and cognitive processes. Only a third of those polled say they have a high level of trust in the accuracy of their data and analytics; one-out-of-five have limited trust for nearly every aspect of the way their organization uses data and analytics.

Aside from concerns, the technology will be "absolutely transformational," says Cliff Justice, principal, innovation & enterprise solutions at KPMG.

"The ability to use technology to augment decisions and effectively automate capabilities that currently require a lot of human labor—from transaction processing in the back office, all the way to call center activities and, on the high end, areas like audit and tax compliance—might be considered the digitation of white-collar research work … You are automating the mundane, repeatable tasks and activities that are best done by machines and technology, refocusing human skill sets and judgment on areas machines can't solve."

While exciting developments are underway, Justice sees the full effect "playing out over a generation."

"You are not looking at one technology that might get outdated and replaced," he says. "These technologies are going to evolve over time. Business and operational models need to change."

Whether they consider AI an evolution of revolution, companies will need to continually assess the quality and quantity of the data they collect and process. "If you consider AI and cognitive technology as a weapon, you need to consider data and content as ammunition," Justice says. "You need to get better at using it and curating it. The best artificial intelligence in the world is going to be useless without expertly curated data to drive and fuel the algorithms."

Among those bringing artificial intelligence to financial firms is Mallinath Sengupta, CEO of NextAngles, a start-up style company operating within Indian technology company Mphasis.

His promise: computers that can process the bulk of 80 percent of routine bank activity so that human compliance officers can focus on the 20 percent that might be a problem. Among the areas his technology focuses on are anti-money laundering alerts, know your customer, financial crimes investigation, liquidity risk management, and keeping pace with regulatory change.

"At current trajectory, rule promulgation in money-laundering, know your customer, and insider trading will shortly render existing compliance teams functionally obsolete," he warns. "There are not funds or competent staff available to keep up. A proliferation of regulations intended to prevent financial crimes and money laundering means institutions are devoting unprecedented resources just to monitor and identify compliance violations. Compliance professionals need assistance in order to retain top talent and focus on material risks rather than data gathering."

NextAngles' concept is an artificially intelligent compliance system that learns from context, not just flagging suspicious patterns but also seeking them out.

"At current trajectory, rule promulgation in money-laundering, know your customer, and insider trading will shortly render existing compliance teams functionally obsolete."

Mallinath Sengupta, CEO, NextAngles

An implementation-level challenge is the collating of data. "Banking systems grew over the past 50 years to be very siloed," Sengupta says. "You need to pull multiple data from different sources to bring them together and get a good view. That is the first pain point."

Rather than unify data systems, he merely wants to be a stopping point for data streams. "Banks have transaction systems, customer profiling, and on-boarding systems," he says. "They already give those feeds to a data warehouse or wherever. All I ask is for them to give me those feeds."

OVERCOMING AI FEARS

The following is from a KPMG client advisory, "Got automatonophobia? Four Steps for Overcoming Your Fear and Getting Started with Process Automation."

Consider the culture of your company, and start with the right-sized pilot.

Your adoption of robotic automation will be based in large part on the culture of your company. Is your enterprise comfortable on the bleeding edge of innovation? Or is it more of a technology follower?

If you are like most organizations, you are not going to rush to cognitive solutions such as IBM Watson out of the gate. Rather, when your culture is on the more risk-averse end of the spectrum, you are wise to build some experience with small, tactical pilots in robotic process automation (RPA) before graduating to more sophisticated, cognitive capabilities.

In finance, for example, if you have employees processing thousands of invoices a day, RPA can drastically reduce costs while improving speed and accuracy. Software robots can be "trained" to extract attributes from invoices in a certain format, enter the invoice data into the enterprise resource planning (ERP) system, progress it through a workflow, and assign it to appropriate approvers.

Create a center of excellence to operationalize automation throughout the enterprise

After you have built some experience with Class 1 automation and want to scale automation at the enterprise level—or if your company is already on the bleeding edge of innovation and is willing to dive right into new enterprise technology—

it is important to create a central automation team. This team, or center of excellence (COE), should consider the role of automation in the enterprise strategy, develop an automation plan, manage implementation, and drive adoption throughout the organization.

Which parts of the business could benefit from automation? What is hype versus reality? How can your company get value from currently available technologies, while also considering fast-emerging technologies? The automation team should answer these types of questions, focusing heavily on strategic alignment and change management.

This team will ultimately take the lead on automation projects throughout the enterprise, including the establishment of standards, response to business units' requests for new solutions, management of vendor relationships, assessment of benefits, and development of controls to prevent rogue "bots" that could cause legal, regulatory, or IT issues.

Identify initial targets for automation projects, and choose the right technologies

While your functional RPA pilots demonstrated the benefits of automation technology, the COE can formally operationalize and scale those solutions, while identifying opportunities for others. To determine enterprise targets for RPA, the COE should identify high-volume, highly transactional process areas.

Assess the benefits of automation projects, and prioritize them accordingly

A key way for the COE to identify project opportunities is by educating the business units on the capabilities of RPA and other classes of automation—and then calling on them to identify automation projects. However, the incoming requests may soon be more than the development teams can handle. That is one reason the COE should continually evaluate the automation benefits and prioritize them accordingly, while staying true to the strategy.

Source: KPMG

AppZen is another startup seeking to marry compliance with artificial intelligence. Among its clients are Hitachi, SunRun and Cantor Fitzgerald, says co-founder and CEO Anant Kale. His focus is a narrow one for now: employee expense reports.

The pitch: due to the high cost of expense auditing, most companies use random sampling and threshold-based auditing to catch policy violations or fraud, which has made them passive and leaves them at risk. With machine-learning algorithms, however, companies can review expense report data, crosscheck it with external sources (including sanctions and Politically Exposed Persons lists) and social media in real-time, and detect compliance problems.

"We saw that managers were really worried about compliance issues and things like the Foreign Corrupt Practices Act," Kale says. "Looking at expenses and trying to find a problem with them is trying to find a needle in a haystack. It takes reading through every document every receipt that is attached, and every justification you are provided. It takes an immense amount of time and trying to understand the context around expenses is a very difficult thing."

His firm's approach is to evaluate and review expenses just as a human auditor would. Every receipt, boarding pass, or whatever documentation there is can be read line by line and be compared to both company policy and regulations, all while being screened against external data sources that range from Yelp and Google to the Treasury Department.

An expense claiming dinner for two at the Olive Garden, for example, will be evaluated to ensure that the claimed cost was within reason and to verify the identity of the recipient. A report from "John's Grille," a lesser known establishment, will be verified. Is it a legitimate restaurant, or a policy-prohibited strip club?

"We can find out if an expense is legitimate or not," Kale says. "It could be somebody's family meals that are being claimed as business expenses, or it could be at an establishment where you really don't want employees to spend money. The most important thing is to really understand the context around attendees. We really need to see if they are legit."

"The idea behind using AI to do this is that now you have your auditors looking through hundreds of documents trying to figure out if there is something wrong," he says. "We filter it down into maybe four or five high-risk areas and present all of this information back to the auditors. The auditors who are doing the job today are ineffective because 90 percent of the time they are not finding anything. Now, they are presented with a subset of data that has already been checked, analyzed, and researched."

Kale stresses the importance of human expertise being the final and most important step in the process.

"The final judgment and decision that has to be made relied on the skills and the training of auditors," he says. "What we are doing then is giving them the chance to apply their skills and knowledge to see if something is non-compliant, and they should be acting on it. That makes them far more effective and solves a compliance problem, because they are doing a lot of work that wouldn't otherwise be looked at."

Automation and IA may also help stem the costs associated with compliance and risk management, especially in financial services.

The increased regulatory scrutiny facing banks has led to massive investments in compliance. For example, HSBC's total expenditure on regulatory compliance in 2015 was $2.9 billion, a 33 percent increase from 2014.

The question: Is arbitrarily boosting compliance personnel effective? The concern: Will investors start to balk?

"One of the things we find with the banks we are working with is that the way they react to regulations now, and will probably continue to in the near term, is that they just hire a ton of people in their compliance department," says Sachin Sachdeva of SSA & Company, a technology consultancy. "They just try to throw bodies at everything. Investors have been ok with it, because they are avoiding risk, but I think the appetite for just throwing money at the problem is decreasing. We need to come up with more creative solutions and using something like natural language processing can have a big impact."

Another issue is that this approach can actually harm compliance efficiency. "We've seen some banks that really staff up their compliance department, but then operational and technology budgets get cut," Sachdeva says. "They have this mismatch where they don't have the staff to implement these policies and are opening themselves up to risk again. That's where automation and using technology as a solution for the lack of manpower can have a big impact."

There are basic operational changes offered by emerging technology that can make the lives of CCOs easier.

"When it comes to taking a regulation, translating it into a policy, and then figuring out how it is going to be implemented operationally, even today a lot of it is very manual and very paper-based, with a lot of PDFs and Word documents going around," Sachdeva says. "The use of natural language processing to automate some of those activities allows humans to review the most important parts of the document, rather than having to review an entire document. Simple things like that can have a pretty big effect." ■

Compliance automation is more than just number crunching

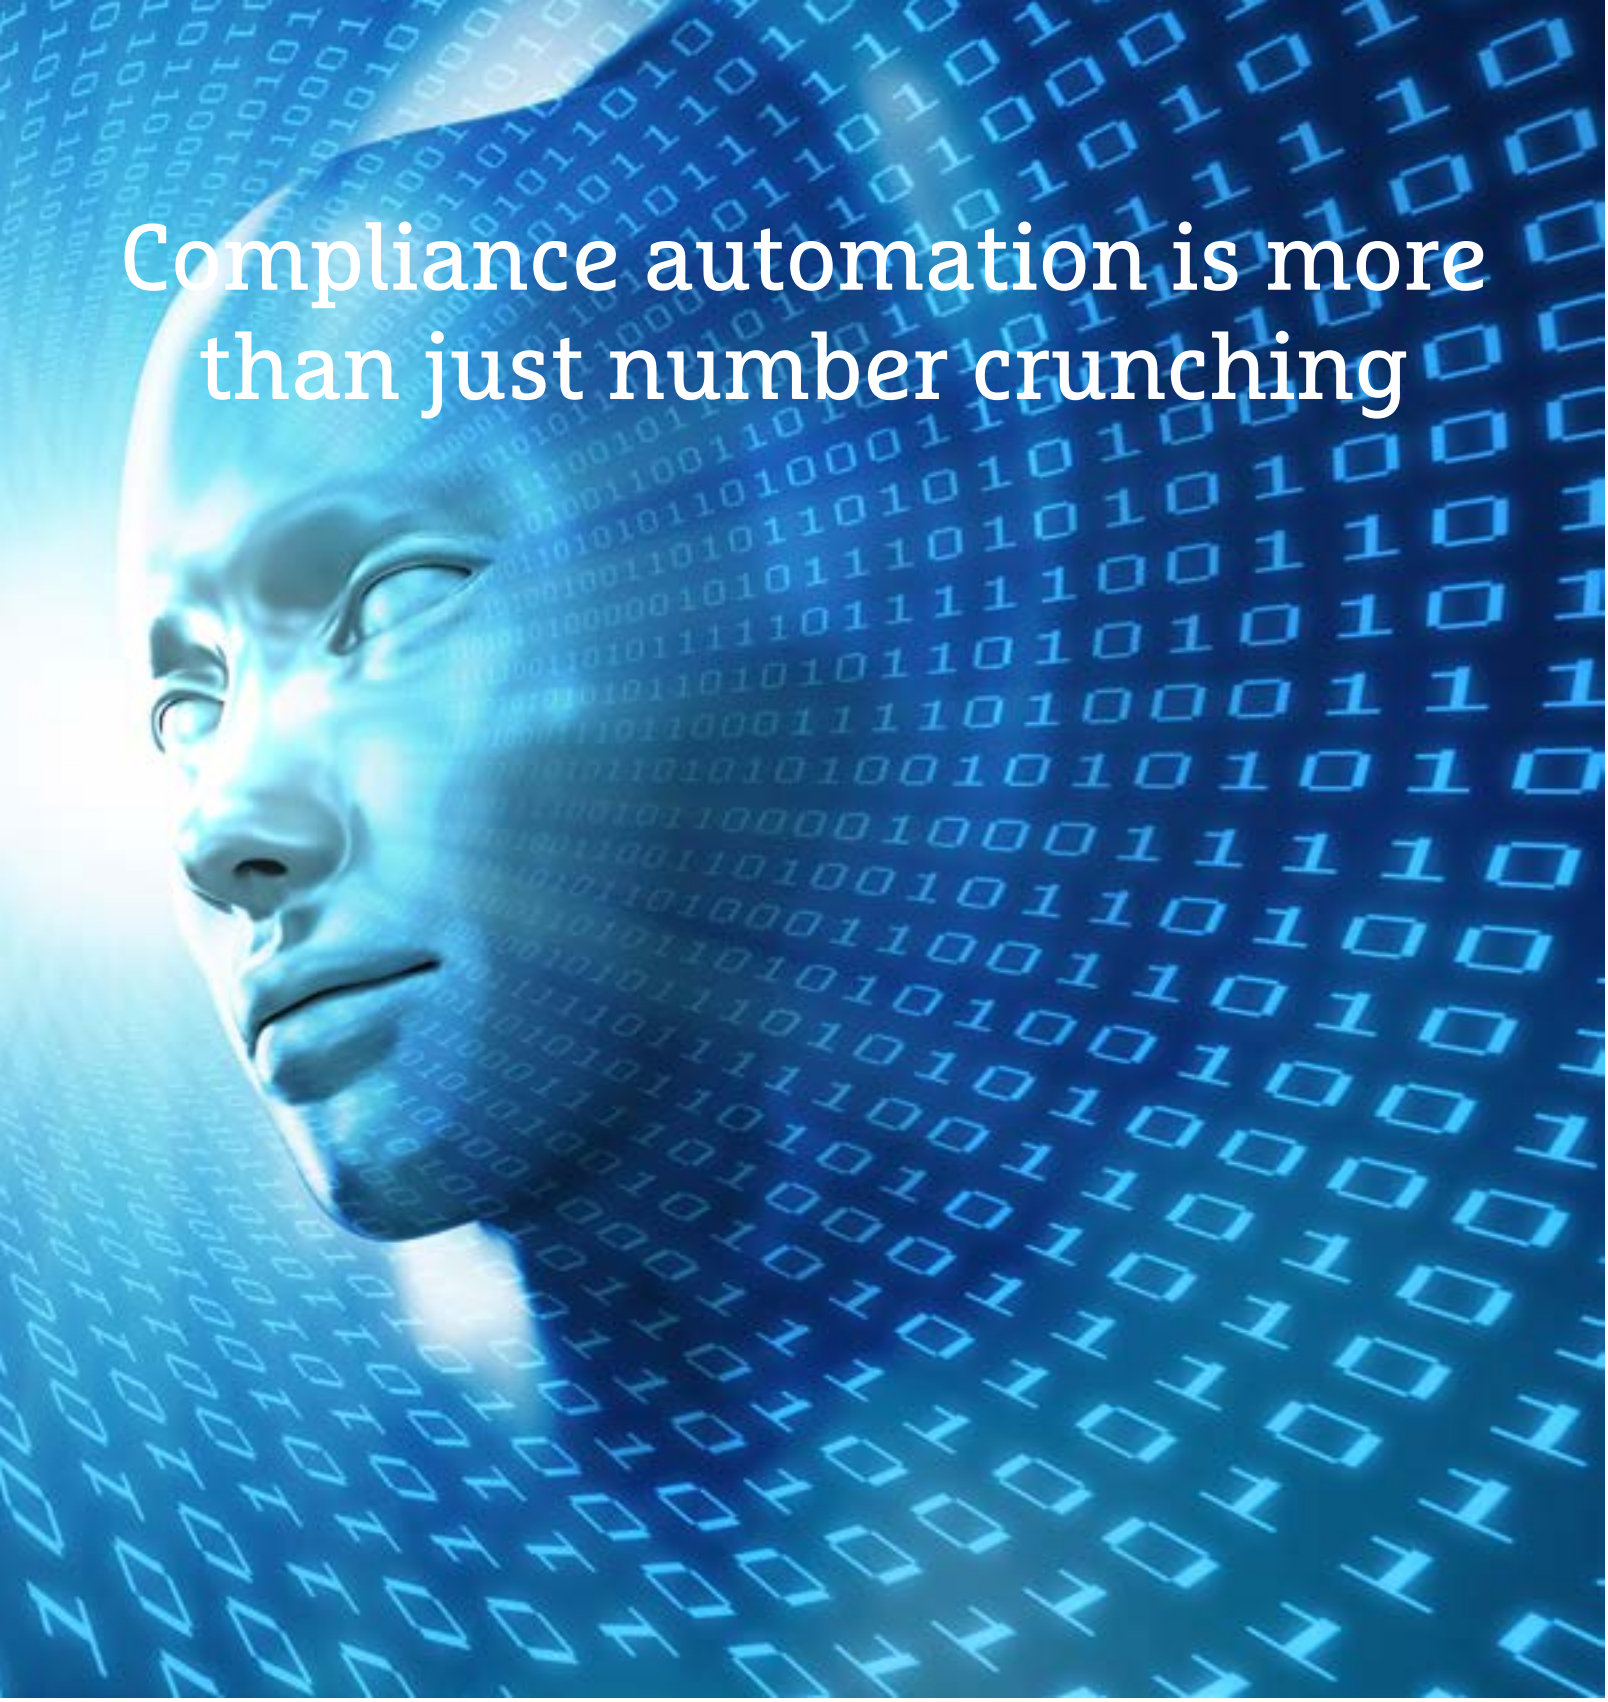

Increasing automation of compliance has much promise, but what can you do with it so far, and how can it improve? **Bill Coffin** reports.

There is a long-running Internet joke that anytime somebody speaks about a potentially troubling advance in artificial intelligence or robotics, the pat response is, "Sarah Connor unavailable for comment." Connor, of course, was the heroine of the Terminator movies and has more reasons that almost any other cinematic character to distrust self-aware machinery.

I think of that joke often when I see the proliferation of stories about AI and robotics. Boston Dynamics recently put out some video of a robot on wheels that can roll around upright and even jump over things, which, if you've ever been phobic about robots, will surely give you nightmares. Sarah Connor unavailable for comment. Likewise, by 2020 we will have computers with enough raw processing power to equal that of the human brain and, by 2050, we'll have computers with raw processing power equaling the sum of every human brain on the planet. Sarah Connor definitely unavailable for comment.

But despite such wise-cracking, there is something very real to the growing use of automation and artificial intelligence within compliance, especially in light of the growing need for measuring effectiveness in compliance programs, and demonstrating a proper return on investment. A significant portion of compliance, especially in the financial services realm, is simply plowing through a large quantity of transactions and communications, and trying to determine if each and every one of them is compliant with every regulation and internal policy that applies. It can be an enormously work-intensive process unless you automate it. Only now is that automation finally bearing fruit. And as it does, compliance professionals are right to wonder what kind of future they might have in a field that is increasingly being given over to thinking machines.

Mallinath Sengupta is CEO of NextAngles, a regulatory software firm that focuses on AML, KYC, and other compliance needs. The way Sengupta puts it, the problem with automation is also in its strength. Computers are very powerful, but very dumb. Without the contextual information that humans take for granted in what we consider to be "common knowledge," no software system can really fully replace human overseers when it comes to compliance. Sure, you can have a system that flags transactions over a certain threshold, but it takes a human to know when those flags make sense—say, a cash-intensive business like a gas station … or when a gas station transacts way more money than is usual for an operation of that type—and when they don't. Case in point: A gas station might regularly log cash receipts of more than $30,000 a day, so the trick is seeing when it suddenly logs more than $100,000 in a day and knowing enough to see that for the red flag that it is. Simply filing everything above a single benchmark isn't good enough.

That is where somebody like Tara Raafat, NextAngle's chief ontologist comes in. Ontology is the science of how concepts connect to each other to form a greater understanding, and that is what Raafat does with data sets and Turing-based machine learning. She designs how machines can draw from existing pools of data to better contextualize certain conclusions being made by particular alerts, which minimizes the need for human analysis. It is a slow process and, at the moment, we can only really enjoy machines eliminating some of the most basic early branches of a fully contextualized decision tree. But it's harnessing the potential of very powerful, very dumb machines to make compliance officers deal with a lot fewer red flags that turn out to be nothing.

Alex Baydin, CEO, of PerformLine—a provider of marketing compliance software—points out that marketing is another area that needs both automation, and the ability to separate wheat from chaff, data-wise. Consumer financial services companies, for instance, need to know every instance where their brand appears on the web, even when the brand managers might not be aware of it. Marketing compliance automation employs crawlers that look for every instance of where a brand appears, understands the content surrounding that brand on a landing page or blog, and runs it against the brand owner's own rules engine to make sure everything is in compliance. This is especially helpful for businesses such as mortgage lenders that have to comply with regulations over truth in lending as well as unfair, deceptive, or abusive acts and practices. Does the brand appear next to content that could be seen as deceptive? Human eyes will never catch everything.

The idea here—and this is something that Sengupta, Raafat and Baydin all stressed—is not to eliminate humans from compliance entirely, but to make sure that the humans who are in compliance can focus on more strategic risk analysis, and get out of the rote fact-checking business. So far, that has been a very human-intensive endeavor that frankly, isn't very sustainable and isn't very cost-effective. Nobody really wants their compliance efforts to ultimately rely on the accuracy of a huge number of people stuck doing a lot of relatively low-brainpower tasks over and over and over every day. Especially not in a Yates Memo era when somebody on cruise control misses a red flag that looks like all the others, and accidentally sets into motion a chain of events that might make headline news.

AI isn't magic or science fiction or the kind of thing you can just buy off the shelf. But what it can do is take a lot of the grunt work out of compliance and free up those resources for other things. As compliance strives to become a faster, smarter discipline, making the right use of the technology at hand is not just best practices, it's becoming common practice.

Sarah Connor might not have a lot to say about that, but that's alright. She never worked in compliance. ■

Jose Taubena
CW Columnist

{ACCOUNTING & AUDITING}

As compliance programs must deal with growing amounts of data, they need tools to help separate true risks from background noise. **Jose Tabuena** explores.

Compliance monitoring and artificial intelligence

With the advent of cyber-security attacks, developers of advanced artificial intelligence security monitoring solutions have emerged. Understanding when and how often monitoring solutions should be executed, however, presents trade-offs to be considered.

Legacy approaches to risk monitoring look for recognized threats by known signatures and pre-built event detection logic. Often these standby methods rest on tech confines and as a result are not aligned to business risk. These limitations can lead to detection challenges such as "data" overload (missing the important needles), "alert" overload (too many false alarms where all the needles look the same), and gaps in skills needed to quickly analyze, recognize, and act on a real event. Trying to monitor every transaction or activity (though essential for some compliance and security functions) to manage threats can be just as ineffective as completely locking down all entry points in an attempt to secure everything.

New approaches to compliance monitoring and threat detection are needed. Successful threat detection and response starts with understanding your top risks, which can be a combination of business, regulatory, and technical risks, including known threats such as data breaches, industrial espionage, fraud, corruption, and disruption of business operations.

The simple pattern-matching approaches of the recent past are highly susceptible to both false positives and false negatives. Advances in machine learning and affordability of large-scale computing resources enable more sophisticated anomaly detection. In order to accomplish this new approach to threat detection, compliance professionals will need greater knowledge of core operational processes to understand a potential compliance incident's business context. In short, the human element remains critical.

Privacy access monitoring for threat detection provides a straightforward example. Open-access environments, in which authenticated employees (i.e., those who have presented valid credentials) can access any patient's record in the system, even if they are not treating that patient) are common across electronic medical record systems, despite their associated privacy risks. The choice to deploy an open-access environment instead of fine-grained access controls is often based on the need for caregivers to access information for continuity of care and in emergencies. For example, if access to a medical record is blocked, the caregiver will not be able to identify the patient's medication allergies and, if given in an acute setting, certain medications may cause harm or even death. Many health firms therefore have traded more granular patient privacy protections for health delivery utility and efficiency.

In an open-access environment, privacy professionals must determine how to best monitor medical record access for inappropriate use. Privacy laws and organizational poli-

cies do not permit curious snooping in the record for those not on the patient treatment team. Manual auditing techniques are difficult to scale to meet the needs of modern healthcare volume, necessitating automated monitoring systems.

The process for reviewing a flagged access involves going through the patient's medical chart to determine if the accessing employee had a clinical or operational reason to do so. This manual process takes time and often result in false positives that result in wasted staff effort and can be overwhelming to a privacy program. Machine-learning systems can leverage operational context to reduce false positives, decreasing the time to complete access reviews.

Health organizations are generally required to log every access to their electronic medical record for years, for security purposes, and to accommodate a patient's right to know who has accessed their record. The challenge is monitoring these large logs. The systems typically record millions of accesses per week, limiting the capability and usefulness of manual auditing approaches. Because of the volume, privacy officials often deploy simple flags to focus on high-risk behavior such as employees accessing records of VIPs, co-workers, patients with the same last name, or family members.

Near real-time monitoring systems are delayed in their ability to alert on suspicious activities but are able to incorporate more clinical context than real-time systems. The addition of context drastically reduces false positive alerts, because the clinical context can be used to filter away accesses that occur for appropriate reasons. Moreover, by auditing for both appropriate and inappropriate accesses, the monitoring coverage drastically increases as the system can analyze more types of access. The ability to automatically audit and filter appropriate access using clinical context can mean the difference between practical management of potential breaches and drowning in alerts.

Reactive monitoring systems have many of the same benefits of near real-time systems, but suffer from long detection delays. Specifically, the system can utilize the complete clinical context to understand and identify suspicious activity, again resulting in broader monitoring coverage than real-time systems. However, breaches may have occurred for months without detection.

When developing a compliance monitoring system, compliance professionals need to bring the right people to the table and ensure that business leaders are actively engaged.

System responsiveness. Real-time monitoring systems are able to react quickly to suspicious activity and can notify compliance staff shortly after an event has occurred. Responsiveness is valuable when the time-to-react is imperative. However, it is important to understand the types of inappropriate activities real-time monitoring systems miss and the mistakes they can produce.

False positive rates. Given the need to respond quickly, real-time systems often look at the activities in isolation, without considering business context. Even if the real-time monitoring system could incorporate all context in its decision-making process, the information may not exist in the system at the time the activity takes place.

Coverage. It is further important to consider the types of activities that real-time systems can detect. If the monitoring system only uses previously specified flags, then the system will not be able to identify other types of inappropriate use.

Filtering appropriate behavior. One of the main benefits of near real-time and reactive monitoring systems is their ability to incorporate context into their decision-making processes and filter away appropriate accesses. The challenge will be defining how to accurately filter away appropriate behavior.

Recent published and peer-reviewed research has developed machine-learning methods to address monitoring challenges. These methods can intuitively filter away appropriate accesses by identifying connections between a patient and the employee accessing the patient's record. Such explanation-based auditing systems can infer relationships from a hospital's data, display them to a privacy officer for approval and, once approved, apply them to future accesses. Using this approach, machine-based learning systems have been shown to filter more than 95 percent of accesses, so staff can focus on truly suspicious behavior.

Keep in mind that tools and technologies are enablers—they are not the foundation of a robust monitoring program. As you move toward the use of compliance intelligence, behavioral analytics, and "Big Data," first ask if more data feeds will lead to more alerts or even more noise, and whether analysts are just going to get buried.

As the use and sharing of data intensifies with a more connected economy (e.g. the Internet of Things), there is an ever-growing need for efficient and workable monitoring approaches. Many organizations can improve the risk alignment of compliance monitoring as they strive to innovate and drive performance—ironically, the very things that magnify compliance risk. ■

Thomas Fox
CW Columnist

{ENFORCEMENT & LITIGATION}

A look from **Tom Fox** at how artificial intelligence is changing the face of compliance in order to incorporate cultural values into the hiring process.

Artificial intelligence meets compliance

Words you do not often see in the same sentence are compliance and artificial intelligence (AI). However, those words came together in a recent article from *Fast Company* magazine around the issue of how "artificial intelligence is being used to screen, test, and hire new talent." One consistent theme from the Justice Department is the integration of technology into a best practices compliance program.

Indeed, in every DPA, Attachment C—which lists out the expectations for a corporate compliance program—is found the following, "The company will conduct periodic reviews and testing of its anti-corruption compliance code, policies, and procedures designed to evaluate and improve their effectiveness in preventing and detecting violations of anti-corruption laws and the company's anticorruption code, policies, and procedures, taking into account relevant developments in the field and evolving international and industry standards."

Companies such as Facebook, IBM, and others are beginning to incorporate AI into their hiring practices. This is moving beyond simply scanning the social media landscape for posts, tweets, and the like. These companies are using algorithms into analyzing facial expressions and word choice during interviews. One developer has come up with a written test to evaluate such soft skills as "grit, curiosity, and polish."

Hiring is not only about getting the best and brightest, but also eliminating those who might engage in illegal or unethical conduct. This could certainly be helpful in the anti-corruption compliance space where even background due diligence checks can fail to note employees who might step over the line and violate such anti-bribery laws as the FCPA.

One consistent theme from the Justice Department is the integration of technology into a best practices compliance program.

The next step from the compliance profession perspective would be to use AI to help incorporate cultural values such as doing business ethically and in compliance into the hiring process. Indeed there are some cutting-edge technologies that look for soft skills. And then it would be to use those same technologies to impress upon the candidate the cultural and ethical value of an organization. AI and compliance may soon be meeting up at an Human Resources department near you. ■

Financial Research Associates

COMPLIANCE WEEK

AI, Technology Innovation & Compliance

Creating the Compliance and Audit Programs of the Future

June 27, 2017 | New York, NY

REGISTER NOW

CALL 704-341-2647

https://events.complianceweek.com/C103

REGISTRATION RATES

In-House Professionals - $495.00

Consultants and Service Providers - $995.00

LIMITED SEATS AVAILABLE AT THESE RATES

- Find out how technology innovation will grow your business
- Get the tools to strengthen your compliance functions
- Hear from early adopters of artificial intelligence
- Watch demos from proven technology innovators
- Learn how to practically apply AI in your organization

#AICompliance