Hedge Fund Law Report
An Acuris company

HFLR.com

September 5, 2019

TECHNOLOGY

AI for Fund Managers: How to Use It to Streamline Operations (Part One of Three)

By Shaw Horton, *Hedge Fund Law Report*

The use of artificial intelligence (AI) in the investment management industry has traditionally focused on quantitative investing. Nevertheless, fund managers can use AI for several back-office functions, including monitoring trader behavior, assisting with the anti-money laundering (AML)/know your customer (KYC) process and improving cybersecurity defenses. Employing AI for these purposes, however, does not come without challenges, particularly for smaller managers.

This article, the first in a three-part series, explores what AI is; how prevalent it is in the funds industry; how it can be used; how fund managers can determine what functions to automate and what obstacles may interfere with implementing AI solutions; and whether humans are still needed in the process. The second article article will analyze what the U.S. government and others are doing to both promote AI and foster its responsible use; how fund managers should diligence and contract with third-party AI service providers; and what risks of bias exist. The third article will evaluate how fund managers can automate their legal departments and what they should do to ensure that they maintain their data subjects' privacy.

For a discussion of machine learning in the quantitative-investing context, see our three-part series: "Dispelling Myths and Misconceptions" (Aug. 9, 2018); "Regulatory Action, Guidance and Risk" (Aug. 23, 2018); and "Special Risks and Considerations" (Sep. 6, 2018).

What Is AI?

The Financial Stability Board (FSB) issued a report on AI and machine learning in financial services in late 2017. The FSB defined AI as the "theory and development of computer systems able to perform tasks that traditionally have required human intelligence."

Machine learning is a subset of AI and is a "method of designing a sequence of actions to solve a problem, known as algorithms, which optimise automatically through experience and with limited or no human intervention," according to the FSB. Machine learning is generally used to predict and categorize, rather than develop causal inferences.

Machine learning can be further sub-categorized as follows:

- Supervised learning involves an algorithm that is "fed a set of 'training' data that contains labels on some portion of the observations," the report states. For example, a manager may feed an algorithm certain contracts and label a subset of those contracts as either containing a source code escrow provision or not. After learning a general rule of classification based on that training set, the algorithm can then classify the remaining unlabeled contracts.
- Unsupervised learning involves an algorithm that is fed a set of data that does not contain labels. "The algorithm is asked to detect patterns in the data by identifying clusters of observations that depend on similar underlying characteristics," the report observes. For instance, a manager may feed an algorithm several contracts without specifying whether they contain source code escrow provisions. The algorithm will then cluster that data – *i.e.*, put them into groups – based on similarities or patterns.
- Reinforcement learning is a middle ground between supervised and unsupervised learning and involves an algorithm that is "fed an unlabeled set of data, chooses an action for each data point, and receives feedback (perhaps from a human) that helps the algorithm learn," explains the report. For example, a manager may feed the algorithm the same set of data as it did under the unsupervised learning example, but the manager then tells the algorithm whether it was correct or incorrect in each of its classifications as to whether a contract has a source code escrow provision.

Deep learning relies on algorithmic "layers" that mimic the "structure and function of the brain." A deep-learning algorithm, which is based on artificial neural networks, can be supervised, unsupervised or reinforced. Deep-learning algorithms have been applied to colorize black and white images; automatically translate text from one language to another; and generate text, among other things.

Prevalence in the Funds Industry

According to KPMG's 2019 Global CEO Outlook – which surveyed CEOs from 11 key industries, including asset management and banking – 16 percent of CEOs have already implemented AI to automate their processes, 31 percent have piloted AI and 53 percent are undertaking "limited implementation." Hedge fund-specific data are sparse.

Nevertheless, Matt Savare, partner at Lowenstein Sandler and member of its technology practice group, opined that most, if not all, quantitative hedge funds are presumably using AI. In addition, he stated that traditional hedge funds are likely relying on third-party software packages with AI embedded in them, regardless of whether they know that to be the case.

"A software licensor may note in its marketing materials that it uses AI to, for example, enhance search results, but managers should not say, 'I have to get Technology X because it has AI in it.' Rather, they should seek out software because it gets the job done," Savare said. "If that software happens to have AI in it, so be it. Managers should be agnostic about the technology used to solve a problem, as long as it solves the problem."

It is likely, however, that fund managers will increasingly adopt AI as computing power continues to increase; hardware costs continue to decrease; cloud services become more prevalent; and data scientists continue to develop more sophisticated solutions in other industries and within financial services. It may not be clear on its face how advances in, for instance, self-driving cars, will be relevant to the financial services industry, but like in mathematics – where solving abstract problems can have future applications in the real world (*e.g.*, number theory, which G. H. Hardy famously said has no practical application and should be appreciated for its beauty alone, later became central to cryptography) - so too may it apply.

In addition, fund managers will be incentivized to use AI for various business and regulatory compliance reasons. Proper use of AI can reduce costs; improve risk management and productivity; augment decision making; and make it easier to comply with regulations, including those pertaining to data reporting, best execution and anti-money laundering. The FSB report also notes that those factors "may also drive 'arms races' in which market participants increasingly find it necessary to keep up with their competitors' adoption of AI and machine learning, including for reputational reasons."

How Can AI Be Used?

"AI is revolutionizing both the front and back end," said Savare. "Many financial services firms, for example, utilize chat bots that carry out conversations with clients. Sometimes the client does not even realize that he or she is taking to a machine. Similarly, firms can utilize AI to extend credit and detect fraud, among other things."

Although not all the above applications may be relevant to hedge funds and other private fund managers, there is still a wide range of potential use cases. "AI, for example, is commonly used to automate the surveillance or monitoring functions of the compliance department," explained Michelle Ann Gitlitz, founder and co-chair of Blank Rome's blockchain technology and digital currencies group, who also has experience advising hedge funds in other emerging technologies, including AI.

More specifically, AI can be used to monitor the behavior (structured data) and communications (unstructured data) of traders to identify suspicious trading activity. The FSB report notes that machine learning that uses natural language processing (NLP) can interpret unstructured data such as email, spoken word communication and instant messaging. Nevertheless, using machine learning in this way may raise issues with respect to a manager's employee surveillance policy.

See our two-part series on mitigating insider trading risks: "Relevant Laws and Regulations; Internal Controls; Restricted Lists; Confidentiality Agreements; Personal Trading; Testing; and Training" (Sep. 27, 2018); and "Expert Networks, Political Intelligence, Meetings With Management, Data Rooms, Information Barriers and Office Sharing" (Oct. 11, 2018).

AI can also be used to analyze laws, regulations and speeches by government officials. For example, the FSB report explains that a manager could use machine learning and NLP to interpret regulations into "common language," conduct an analysis of that converted text and, finally, "codify the rules for automation into [its] integrated risk and reporting systems" to heighten compliance. Gitlitz noted that machine learning and NLP can likewise extract rules from other documents, facilitating managers' tracking, analysis, creation and processing of legal documents.

"The majority of our productionized AI services relate to extracting insights and predictions from unstructured natural language," noted Michael Gallina, CEO and founder of Phylum Data Suite, a technology firm that offers AI products to the asset management industry and others. "Extracting and modeling around legal documents is not yet commodified, but it is a constant recurring theme in our interactions with hedge funds." The third article in this series will contain more detailed discussion of how managers can implement AI into their legal departments.

"Assistance with the AML/KYC process was one of the first use cases of AI in the asset management industry," explained Gitlitz. Although hedge funds are not required to adopt AML programs under U.S. Department of the Treasury (Treasury) Financial Crimes Enforcement Network rules, other parties – such as banks – may require hedge funds to implement AML policies. Other jurisdictions, including the Cayman Islands, also impose AML obligations on private funds. The AML/KYC process is often time-consuming and expensive, but machine learning can, according to the FSB report, "perform[ ] identity and background pre-checks" by "(1) evaluating whether images in identifying documents match one another, and (2) calculating risk scores . . . [to] determine which individuals or applications need to receive additional scrutiny." A manager can also use machine learning to conduct ongoing checks to revise risk scores, for instance, by accessing public records or social media.

Lex Sokolin – global financial technology co-head of ConsenSys, who has assisted hedge funds and others with their strategic decisions regarding, among other technologies, AI – remarked that AI-based cyber-threat management has gathered traction within the industry. "One approach involves AI that scans the entire footprint of the web and monitors security-level information about all entities that can be tracked. This approach relies on big data web crawling and AI-led structuring of that data in real time," he clarified. "A second approach involves software that is written by AI in order to respond to the millions of new bugs and viruses that are continually being launched. As a new virus evolves, AI software can generate bespoke wrappers like a cyber immune system."

AI can also help managers with the following:

- *Machine-Executable Regulatory Reporting:* "[S]ubstantial errors, blank fields, and other data-quality issues are often more prevalent in new datasets," the FCB report states. Machine learning algorithms can detect those anomalies, which can lead to higher-quality reporting.
- *Risk-Management Processes*, by identifying potential risks before violations occur: The FCB report warns, however, that algorithms may miss new risks or events if they overtrain on past events. "It is also unclear whether a manager needs AI or data science – *i.e.*, the use of statistical learning as opposed to machine learning – for this," said Sokolin.

- *Decision Making*: For example, IBM has developed Project Debater, which is an "AI system that can debate humans on complex topics . . . to help people build persuasive arguments and make well-informed decisions."[1]
- *Extraction and Normalization of Alternative Data*: "Data sets extracted from the web or other unstructured sources can be literally anything," Gallina explained. "It is entirely dependent on the needs of the client."
- *Asset Valuation*, relying on historical data, economic indicators, growth predictions and other data. See "Three Approaches to Valuing Fund Assets and How Auditors Review Those Valuations" (May 11, 2017).
- *Best Execution*. See "How to Avoid the Eight Best Execution Compliance Issues in OCIE's Latest Risk Alert" (Aug. 30, 2018)

Determining What to Automate, and Roadblocks to Adoption

What to Automate

"The 'best' processes to automate – *i.e.*, those that are easiest to time to market – are those that have very clear and fixed parameters for training a system to evolve in a streamlined way," said Gallina. "The greater the potential for extraneous factors, the more difficult it is to either create a rule-based automation routine or a learning-based routine. In either case, anomalous factors must be continually added to an automation algorithm."

Certain back-office processes, like risk management and valuation, tend to be very dynamic, Gallina continued, which means that they may not be the easiest to tackle. One example of a process that he and his team implemented that was not "easy" (but that did have consistent factors) involved extracting problematic issues from logs. "Although issues had to be extracted from natural language, there was a fairly small set of vernacular that applied to the types of problems we were tracking. Additionally, the logs had a very organized schema with regard to how data was inputted," he said. "In other words, it was easier to train a model because the amount of variable ways people logged problems was narrow."

"AI works best for activities that are paper-intensive and repetitive – *i.e.*, high-volume work that is prone to human error," added Gitlitz. "With that being said, before determining what to automate, a firm should think about its goals and what value it seeks to get from automation."

Potential Roadblocks

"It costs a lot to onboard these processes, so managers – particularly those on the smaller end of the spectrum – must carefully weigh the costs and benefits of using AI," Gitlitz continued. "Everyone should be involved – whether that is the technology leader or the compliance leader – when thinking about these issues."

In addition, there remain other roadblocks to AI's adoption. Leadership, for instance, may be hesitant to trust new technologies. Managers and investors may also have ethical concerns surrounding bias, model opacity, quality of data and data privacy.

"A fairly typical challenge in the industry flows from the need for secrecy. We are often asked to solve a problem without knowing the full scope of that problem, or we are asked to look at data without knowing the goal," said Gallina. "Building good machine learning models is about underlying data, the quality of that data and the relationships within that data. Thus, it is difficult when a firm seeks to have its problems compartmentalized."

Sokolin added that "many AI companies get bogged down in selling services on premise to financial institutions, who have to protect their data and are unable at times to participate in modern cloud services due to legacy infrastructure."

"The evolving and very specific needs of hedge funds create another challenge," Gallina continued. "Each scenario requires some degree of customization. As a result, it can be rather challenging to offer turnkey software-as-a-service solutions."

By way of example, Gallina discussed how Phylum deals with NLP. "We have fairly generic models for market intelligence that extract things like 'who' (*e.g.*, who is using a product?), 'how' (*e.g.*, how is an individual using a product?) and 'where' (*i.e.*, physical locations, not geography)." Once clients begin using its core models, they begin to want to customize them to identify a specific problem or data point they are evaluating. "For example, a firm may want further segmentation on the 'who' to identify gender from article comments where the commenter's gender is not apparent. This also dips into potential ethics issues as well," he explained.

Implementation Challenges for Smaller Managers

"Large managers – particularly quant managers – hire armies of data scientists to manage their data and improve their algorithms. Smaller managers simply do not have the same in-house technical capabilities," said Savare. "They have core teams, which may include three or four people, and outsource a lot of their information technology functions to put their systems in place." In those cases, he recommended that a manager conduct triage and identify critical systems, including the back-office applications that are needed to run day-to-day operations, and then seek out areas in which AI could be a good complement.

"Small funds should rely on the expertise of their outsourced technology teams," continued Savare. "Vendors can advise funds on what is needed and what is nice to have. For instance, a fully featured software package, which may cost ten times more than a streamlined version, may not be suitable at the onset. Instead, the manager can upgrade as it grows." The challenge for small hedge funds, he said, is "resources, resources and resources."

"Smaller firms have to rely on their large software, custody or administration providers for productivity tools," explained Sokolin. "That only gets a firm so far, however. To the extent possible, a manager should hire somebody who has a strong mathematics and computer science background to make sense of whether the AI solution adds value or is mere marketing speak." He added, "A manager could also use outside industry consultants to help with vendor management – that may get it to the level of best practice, but not differentiated thought leadership."

Human Element Required? Fostering Investor Trust

Scott W. Bauguess, Deputy Chief Economist and Deputy Director in the Division of Economic and Risk Analysis, noted in a 2017 speech that "the most advanced machine learning technologies used today can mimic human behavior in unprecedented ways, but higher-level reasoning by machines remains an elusive hope." Thus, he predicted that "human expertise and evaluations [will always] be required to make use of the information" provided by technology.

In a 2016 speech, Bauguess stated that:

> machine-learning initiatives require care and thought on the back end, from the human user, who needs to bring [his or her] experience to interpreting the results and deriving meaning beyond simple (or not so simple) correlations. In all cases, the human role is essential when using the results to inform on critical decisions related to policy issues or risk assessment.

Like Bauguess, former CFTC Chairman J. Christopher Giancarlo argued in a 2018 keynote speech that AI and machine learning are meant to assist human efforts by, for example, freeing "staff from repetitive and low value tasks to focus on high value activities that require . . . expert judgment and domain knowledge."

Gallina added that "humans are definitely needed to validate a particular model in its earliest stages and, in most cases, afterward as well to tune the model for best results." In a perfect scenario, he continued, once a desired result of accuracy is achieved and maintained, it is feasible to apply some form of automated checks and balances, allowing humans to step away from the process. Nevertheless, Gallina opined that there are certain areas where humans should always be involved, including moderating bias, given that "every business has an ethical responsibility to the community. It is, of course, difficult to impose specific regulations or rules across the board to enforce this."

When asked about whether investors may be averse to managers who use AI given potential ethical concerns, Savare indicated that a manager should be in good standing as long as it has the appropriate internal controls in place. "Registered investment advisers have a fiduciary duty to their clients, which means that there should be checks and balances when using AI or any kind of algorithm," he said. "Companies like Betterment, which is very highly automated, are doing very well. So, there are clearly a lot of people for whom a heavy reliance on AI is not an issue."

In addition to having robust controls, Savare added that, as long as AI is supplementing intelligent individuals who have experience and a track record of beating the general market, its use will likely be viewed positively.

For more on robo-advisers, see "SEC Settles First Two Enforcement Actions Against Robo-Advisers" (Feb. 14, 2019); and "What Robo-Advisers Can Expect From SEC Examinations" (Jun. 21, 2018).

[1] See https://www.research.ibm.com/artificial-intelligence/project-debater/.

September 12, 2019

TECHNOLOGY

AI for Fund Managers: Government Guidance, Service-Provider Negotiations and Risks of Bias (Part Two of Three)

By Shaw Horton, *Hedge Fund Law Report*

Under the Trump administration, the U.S. government has prioritized American innovation in artificial intelligence (AI). Although the SEC and CFTC have said little about AI, other domestic and international governmental agencies have issued guidance on how firms can responsibly use AI, as well as diligence and contract with service providers who may provide AI solutions.

This article, the second in a three-part series, analyzes what the U.S. government and others are doing to both promote AI and foster its responsible use; how fund managers should diligence and contract with third-party AI service providers; and what risks of bias exist. The first article explored what AI is; how prevalent it is in the funds industry; how it can be used; how fund managers can determine what to automate and what obstacles may interfere with implementing AI solutions; and whether humans are still needed in the process. The third article will evaluate how fund managers can automate their legal departments and what they should do to ensure that they maintain their data subjects' privacy.

See our three-part series on big data: "Its Acquisition and Proper Use" (Jan. 11, 2018); "MNPI, Web Scraping and Data Quality" (Jan. 18, 2018); and "Privacy Concerns, Third Parties and Drones" (Jan. 25, 2018).

President Trump's Executive Order on AI

On February 11, 2019, President Trump issued an executive order on maintaining American leadership in AI.

"AI is critical not only to the economy and society, but also to national security," said Matt Savare, partner at Lowenstein Sandler and member of its technology practice group. "Therefore, it is not surprising that the federal government is taking it so seriously."

The executive order states that the so-called "American AI Initiative" is guided by five principles, including that the U.S.:

- develop "appropriate" technical standards and reduce barriers to entry to facilitate the creation of new AI-related industries and AI's adoption;

- train American workers so that they can "develop and apply AI technologies" in the present and the future;
- engender "public trust and confidence in AI . . . and protect civil liberties [and] privacy"; and
- promote an international environment that supports and is open to American AI and protects related intellectual property.

The executive order further states that the heads of implementing agencies should consider AI as a research and development priority; "take this priority into account when developing budget proposals"; and collaborate with the private sector as appropriate.

For more on the Trump administration, see "How the Tax Cuts and Jobs Act Will Affect Private Fund Managers and Investors" (Feb. 22, 2018); "OCIE Associate Director Outlines Coordinated Compliance Effort Under Trump Administration" (Oct. 19, 2017); and "Pro-Business Environment of New Administration Continues to Have Challenges and Pitfalls for Private Funds" (Sep. 14, 2017).

Treasury Department Report on Nonbank Financials, FinTech and Innovation

In July 2018, the Treasury issued its report to President Trump on nonbank financials, financial technology (FinTech) and innovation, pursuant to Executive Order 13772. The report recommends that, given the benefits of AI and machine learning, regulators refrain from imposing "unnecessary burdens or obstacles" and provide clarity that enables testing and responsible deployment.

In addition, the report suggests that regulators:

- engage with the Select Committee on Artificial Intelligence and collaborate with other agencies;
- appropriately emphasize "human primacy in decision making for higher-value use-cases relative to lower-value use-cases";
- consider model transparency;
- scrutinize a model's robustness against manipulation; and
- hold human beings accountable.

SEC

The SEC has issued no official guidance on the use of AI or machine learning by fund managers, outside of the context of quantitative investment models. Even there, the literature is sparse.

The SEC's Strategic Hub for Innovation and Financial Technology portal links to four speeches and presentations by Scott W. Bauguess, Deputy Chief Economist and Deputy Director in the Division of Economic and Risk Analysis. In a 2018 speech, Bauguess noted that while neither Regulatory Technology (RegTech) nor Supervisory Technology have reached maturity, "both offer significant promise by way of improved market functioning and increased operational efficiencies." The SEC itself, he stated, uses AI to detect market misconduct, relying on many open-source methods. Bauguess's 2017 and 2016 speeches delineated the SEC's efforts in greater detail, including its use of natural language processing, and the role of big data in AI. In a 2015 presentation, he discussed, for example, how the SEC conducts its hedge fund risk assessment.

See our three-part series on open-source software: "What Is Open-Source Software, and How Are Fund Managers Using It?" (Feb. 21, 2019); "What Are the Benefits and Risks of Using Open-Source Software?" (Feb. 28, 2019); and "How Fund Managers Can Mitigate the Risks of Open-Source Software" (Mar. 7, 2019).

CFTC

Unsurprisingly, the CFTC has also turned its attention to AI and other FinTech solutions. In a November 2018 keynote address at Georgetown University Law School, former CFTC Chairman J. Christopher Giancarlo noted that the CFTC is addressing the "digitization of modern markets" through its Technology Advisory Committee, its market intelligence branch and LabCFTC, which seeks to promote responsible FinTech innovation in the markets and increase the CFTC's own adoption of those technologies.

The analysis of increasing amounts of data, which is driven by AI and machine learning, is the most important area on which the CFTC must focus, Giancarlo said. Indeed, he argued that AI is "likely to become as ubiquitous to our commodity and financial derivatives markets as the Internet has become today. This means that all organizations and actors – including market regulators like the CFTC – will need to keep pace with the advance of AI in order to succeed."

Giancarlo expressed optimism that AI and machine learning can be used by the CFTC and other regulators to enhance oversight and enforcement, and to digitize rulebooks, which could lead to dynamic rules and regulations.

For more from Giancarlo, see "CFTC Chair Calls for Reset on Cross-Border Swaps Regulation" (Nov. 29, 2018); and "New CFTC Chair Outlines Enforcement Priorities and Approaches to FinTech, Cybersecurity and Swaps Reform" (Nov. 9, 2017).

FINRA's Innovation Outreach Initiative

Like the SEC and CFTC, FINRA also uses AI – specifically machine learning – to monitor the markets, given that it must process and analyze "roughly 60 billion market events" every day to "find instances of abusive activity."

FINRA also actively monitors FinTech in the securities industry through its Innovation Outreach Initiative, which has resulted in regional roundtables and conferences that have focused, in part, on AI.

See our series on FINRA's 2019 RegTech Conference: "AI and Big Data; Blockchain; and Regulators' Views" (Mar. 21, 2019); "Current Uses of RegTech and Considerations Before Deployment" (Mar. 28, 2019); "Digital Identification, Suspicious Activity Reporting and Machine Learning" (May 16, 2019); and "AI, RegTech Adoption and Compliance Challenges" (May 30, 2019).

As part of that initiative, FINRA requested comments from the industry in 2018 on how it can support FinTech development consistent with its mission of investor protection and market integrity, particularly with respect to AI. The notice states that, although the "use of [AI] in the securities industry has the potential to improve operational effectiveness and efficiency," it also raises unique issues for compliance, including with respect to FINRA Rule 3110, which "requires a firm to establish and maintain a system to supervise the activities of its associated persons that is reasonably designed to achieve compliance with the applicable securities laws and regulations and FINRA rules."

In a September 2018 report on RegTech in the securities industry (FINRA Report), FINRA suggested that firms revisit their written supervisory procedures and "appropriately update[ ] and test[them] to reflect any required changes . . . due to the integration or adoption of new RegTech tools." This may include:

- creating a cross-disciplinary group that is involved in the development, testing and implementation of AI. "Testing of various scenarios and outputs . . . with input from a cross-functional group may . . . help limit potential issues," argues the FINRA Report;
- maintaining simple summaries that describe the AI so that non-technical staff "are better able to assess results that do not align with expectations";
- developing a data quality risk management program to "ensure accuracy, completeness, and consistency of . . . data"; and
- establishing policies and procedures that help a firm "identify, respond to, and mitigate material risks that may manifest in the event errors or malfunctions arise," including adopting alternative processes if the AI fails.

European Commission Guidelines

In April 2019, the European Commission (EC) published ethical guidelines for trustworthy AI, stating that trustworthy AI must be lawful, ethical and robust. To meet that goal, the guidelines set forth seven requirements for AI systems:

1. *Human Agency and Oversight*: According to the guidelines, "AI systems should support individuals in making better, more informed choices in accordance with their goals." To achieve this, firms must properly oversee the systems, taking note that less oversight generally requires "more extensive testing and stricter governance."
2. *Technical Robustness and Safety*: AI systems should be secure against vulnerabilities; have a contingency plan in the event something does go wrong; be accurate; and have reproducible and reliable results.
3. *Privacy and Data Governance*: AI systems should preserve data subjects' privacy (both internally, by restricting access, and externally) and be fed high-quality data (both during training and operation).
4. *Transparency*: Firms should document the data sets, processes and algorithms used by AI, as well as the decisions those systems make. To the extent possible, firms should be able to explain the AI system's decision-making process and the rationale for its deployment, among other things.

5. *Diversity, Non-Discrimination and Fairness*: Firms should remove discriminatory data sets and oversee the development of the AI to counter the introduction of programmer bias. Additionally, to the extent applicable, firms should solicit feedback from those directly or indirectly affected by the system.
6. *Societal and Environmental Wellbeing*.
7. *Accountability*: Firms should conduct internal and external audits of the AI system, as applicable; protect those who express legitimate concerns about the system; and provide remedies if "unjust adverse impact occurs."

The EC suggests several technical and non-technical methods to satisfy these requirements, including procedures or constraints on procedures anchored in the AI system's architecture; a robust testing and validation process; an updated code of conduct; and training of and involvement by stakeholders.

Third-Party Service Providers

Before even approaching a vendor, a manager must identify the problem; evaluate its options; and determine whether the proposed solution is viable and affordable. Only then will the other considerations apply.

"When a manager uses AI to make decisions, the manager is still responsible for those decisions. Of course, the negative consequences could be just as harsh if the manager did not use AI and its manual processes failed," said Michelle Ann Gitlitz, founder and co-chair of Blank Rome's blockchain technology and digital currencies group, who also has experience advising hedge funds in other emerging technologies, including AI. "Nevertheless, a manager must consider its liability before it implements AI. This includes thinking about which risks AI could mitigate, as well as which risks AI could amplify or create." She added, however, that a manager can contract with an AI service provider to assign responsibility. "A manager must carefully review its contracts and understand what falls on the vendor and what falls on it, whether it be in tort or by statue."

Gitlitz recommended that a manager also ask prospective service providers about the standards that the AI system must meet; about the discriminatory decisions the AI could make; about the provider's compliance with relevant privacy laws and other regulations; and whether the system will make the fund's systems less secure – *i.e.*, more vulnerable to a cyber attack.

For example, the FINRA Report notes that new technologies "may involve linking to and pulling in data from multiple internal and external sources on an ongoing basis, which could potentially lead to new sources of security risks." In addition, a manager may expose itself to new or additional risks as it provides vendors access to its systems.

Managers that are concerned with explainability may also choose to only select machine learning models that follow supervised learning, advised Lex Sokolin, global financial technology co-head at ConsenSys who has assisted hedge funds and others with their strategic decisions regarding, among other technologies, AI. He explained that the training of a supervised learning model is a very clear process: "You know what motivates its decision making, what factors it weights (often) and how it got there. With unsupervised learning, however, the weights and factors are not clear to a human being – they are derived mathematically but are not necessarily translatable." He added that "explainability is not as fully baked as some of the providers may want to claim."

"If an AI company offers its solution as a software-as-a-service, it may have access to a manager's systems, nonpublic financial information and other data. In those cases, conducting diligence with respect to, and including contractual provisions regarding, data security is very important," said Savare. "On the other hand, if the AI solution is installed on premises, this is less of a concern because the company will likely not have the same access to the manager's data."

Firms should also reference the Information Technology Examination Handbook on Outsourcing Technology Services (OTS Booklet) published by the Federal Financial Institutions Examination Council (FFIEC). The FFIEC is a formal interagency body that prescribes "uniform principles, standards, and report forms for the federal examination of financial institutions" and makes "recommendations to promote uniformity in the supervision of financial institutions."[1]

Although the FFIEC is composed of banking regulators[2] and the OTS Booklet does not specifically mention AI, the recommendations in the OTS Booklet can serve as a useful guide for fund managers when soliciting and employing third-party service providers for AI solutions. In particular, properly assessing service providers can mitigate reputational, strategic and compliance risks.

The OTS Booklet recommends that a firm:

- evaluate the "quantity" of risk associated with the function outsourced (*e.g.*, by asking about the sensitivity, criticality and volume of data), the service provider (*e.g.*, by asking about its financial condition, turnover, experience and reliance on subcontractors) and the technology (*e.g.*, by asking about its reliability, security and scalability);
- define its "business requirements" by documenting its expectations of the outsourced services, including with respect to scope and nature; standards; minimum acceptable service provider characteristics; monitoring; duration; and protections against liability;
- generate requests for proposal (RFPs) based on the information developed in the business requirements phase. According to the FFIEC, an RFP should "describe the institution's objectives; the scope and nature of the work to be performed; the expected production service levels, delivery timelines, measurement requirements, and control measures; and the [firm's] policies for security, business continuity, and change control." Prospective service providers should address each requirement, and managers should evaluate those answers against their needs;
    - conduct due diligence on prospective service providers, including on:
    - corporate history;
    - qualifications, backgrounds and reputations of principals;
    - financial, operational and technical soundness;
    - reputation;
    - bandwidth and effectiveness of services;

  - internal controls and security history for safeguarding of data;
  - legal and regulatory compliance, and the existence of any litigation or regulatory actions. Note that the Financial Stability Board report on AI and machine learning in financial services warns that AI service providers "may fall outside the regulatory perimeter or may not be familiar with applicable law and regulation," which increases the risk that they "may not be subject to supervision and oversight";
  - insurance coverage;
  - disaster recovery and continuity plans; and
  - management style and culture;
- negotiate a contract with the chosen service provider(s). Contracts should, among other things:
  - define the scope of service, including required activities; time frame for implementation; how the service provider's product will interact with preexisting systems; and whether the service provider will provide software support, maintenance, training or other services;
  - outline performance standards and remedies for failure to meet those standards;
  - establish confidentiality and security standards prohibiting the service provider from using or sharing firm or client data with others (except, for example, to the extent necessary to provide services) and requiring it to disclose breaches and corrective actions taken;
  - address controls regarding, among other things, compliance with regulations, record maintenance, access to records and notifications/approval rights for material changes;
  - specify audit rights, including frequency and whether audits will be conducted internally or externally, as well as what other reports will be provided to the firm;
  - stipulate how the service provider will back up and protect records, and test and implement disaster recovery plans;
  - describe the ownership and permissible use of the firm's data, hardware, software and other intellectual property rights; and
  - detail the duration of the contract, taking into account the rapid change of technology and notice periods required for non-renewal; and
  - continually monitor the service provider. The oversight program should, according to the FFIEC, "monitor the service provider environment including its security controls, financial strength, and the impact of any external events."

A manager should also create a contingency plan in the event that the service provider is no longer able to fulfill its obligations, particularly if AI and machine learning are used for "mission-critical" operations.

Ethical Considerations: Bias

"Bias is certainly a major consideration for other financial services companies," said Savare. "A credit card company, for example, may use AI to determine whether to extend credit to those without credit histories. Without traditional data points, they must rely on other metrics to do so." He continued, "If the AI has certain biases – for example, the programmer may have had an implicit bias when coding the algorithm – it may negatively affect certain groups, all things being equal."

Employment decisions are another area into which bias may creep, and one more relevant to hedge funds and other private fund managers. The AI may screen employees but do so unfairly based on gender, age or race. "The software makers who are coding their systems, or the employers who are utilizing the AI and feeding more data into it, must be mindful that the AI's decisions be fair and equitable," Savare cautioned.

"AI systems are not themselves biased. They are mirror reflections of what the underlying data on which they are trained displays from past human behavior and decision making," explained Sokolin. "Therefore, the real audit should be around what training data sets were used in order to stand up the systems, as well as the automated decision making that they are enabled to generate."

See our four-part series on diversity: "Why Equal Representation Within Fund Managers Is Essential" (Oct. 4, 2018); "Ways Fund Managers Can Promote Diversity and Inclusion" (Oct. 11, 2018); "What Implicit Biases Are and Whether Interventions Are Effective" (Oct. 18, 2018); and "How Constrained Decision Making, Along With Legal and Compliance Leadership, Can Help Reduce Fund Manager Bias" (Nov. 1, 2018).

Malicious actors may also introduce bias into algorithms intentionally. For example, competitors could release biased data to negatively affect those training their algorithms. Additionally, hackers could deliberately introduce bias, which underscores the need for a robust cybersecurity program.

See our three-part series on how fund managers should structure their cybersecurity programs: "Background and Best Practices" (Mar. 22, 2018); "CISO Hiring, Governance Structures and the Role of the CCO" (Apr. 5, 2018); and "Stakeholder Communication, Outsourcing, Co-Sourcing and Managing Third Parties" (Apr. 12, 2018).

Regardless of whether the bias is intentional or unintentional, it can lead to reputational or financial harm.

[1] See https://www.ffiec.gov/about.htm.

[2] That is, the Board of Governors of the Federal Reserve System, the Federal Deposit Insurance Corporation, the National Credit Union Administration, the Office of the Comptroller of the Currency and the Consumer Financial Protection Bureau.

Hedge Fund Law Report
An Acuris company

HFLR.com

September 19, 2019

TECHNOLOGY

AI for Fund Managers: Automating the Legal Department and Maintaining Privacy (Part Three of Three)

By Shaw Horton, *Hedge Fund Law Report*

Artificial intelligence (AI) can be used to increase the effectiveness of in-house legal departments, allowing attorneys to focus more of their time on value-adding, strategic services. Leveraging AI in this way does, however, entail risks, so general counsels (GCs) must ensure it is properly integrated. Fund managers must also reassess their compliance with applicable privacy laws given that AI may increase the types of data that they collect.

This article, the third in a three-part series, evaluates how fund managers can automate their legal departments and what they should do to ensure that they maintain their data subjects' privacy. The first article explored what AI is; how prevalent it is in the funds industry; how it can be used; how fund managers can determine what functions to automate and what obstacles may interfere with implementing AI solutions; and whether humans are still needed in the process. The second article analyzed what the U.S. government and others are doing to both promote AI and foster its responsible use; how fund managers should diligence and contract with third-party AI service providers; and what risks of bias exist.

See "How Fund Managers Can Use Technology to Transform and Streamline Complex Legal Operations: One Manager's Example" (Jul. 18, 2019).

Automating the Legal Department

Using AI to Conduct Routine Legal Work

"It is important to distinguish between routine and non-routine legal work," said Matt Savare, partner at Lowenstein Sandler and member of its technology practice group. "AI is currently at a stage in its lifecycle where delegating routine legal tasks can yield enormous benefits." For instance, due diligence is generally a routine function that is often done by first- or second-year associates or contract attorneys.

Savare illustrated this by offering the example of a large quantitative hedge fund, which had hundreds of different data-licensing contracts whose documentation resided in many different forms, ranging from clickthrough agreements to PDFs executed years prior. The fund sought to find the most cost-effective way to parse through all the licensing agreements and extract the most important provisions, such as how and how long the data can be used, and whether there were any territorial restrictions.

"Although they were only concerned with a discrete number of provisions – perhaps ten terms – they wanted us to look through hundreds or thousands of individual contracts," he recalled. "I told them that it did not make sense to use an army of contract attorneys. Instead, it would be better to engage an AI service provider that could use optical character recognition, or OCR, to convert non-digital data into machine-encoded text and then flow it through its AI platform." He continued, "The AI platform could then conduct a large portion of the diligence for them, including summarizing contract provisions, finding material information and identifying anomalies."

AI works well when it ingests large amounts of data. "To create better algorithms, you must feed the machine," Savare said. "Because these types of contracts are very similar in structure, it is easy for AI to pull out the relevant elements of the contracts into a digestible report." At that point, attorneys can review the report and conduct spot checks for quality control. "If the machine analyzed things correctly, we tell it that. If it did things incorrectly, we tell it how the task should have been done, and it learns," he explained.

The same process can be applied when a hedge fund or private equity fund is looking to invest in or buy a company, as applicable, where armies of corporate attorneys would traditionally comb through tens of thousands of pages to see whether, among many other things, a contract is assignable or there are issues related to a change of control.

AI can also be used effectively in research. In the past, if an attorney wanted to find a case that stood for a certain proposition, he or she would conduct a search on LexisNexis or Westlaw, which would return dozens, if not hundreds, of results and be reliant on that attorney's ability to determine appropriate keywords. The attorney would then need to analyze whether those cases are relevant, persuasive or controlling. "AI, on the other hand – particularly AI that uses natural language processing – can do that first level of review," said Savare.

Freeing Up Lawyers to Focus on Value-Add Services

"There is only so much time in the day. If an attorney can focus more of his or her time and energy on high-value, strategic services, then he or she will have more meaningful touch points throughout the day," opined Savare. "I am not suggesting that attorneys eschew the fundamental tasks altogether but rather that they reprioritize by allowing AI to assist with those tasks.

AI, Savare continued, can help an organization service more clients; be more organized and consistent; use fewer people; and develop deeper connections among employees.

Limitations

Ultimately, AI has not displaced partner-level involvement. "You still need experienced attorneys to, among other things, provide strategy; create research parameters; conduct quality control; and impart wisdom based on their knowledge of the parties, industry customs and prior dealings," explained Savare. "AI is not a panacea; human involvement is still needed."

"AI is only as good as the programming that went into it – there may be bias, or it could have been coded incorrectly." Savare continued, "Garbage in means garbage out." Thus, investment managers and other fiduciaries cannot outsource everything in this way because they are ultimately in charge of providing the service. "If the AI makes a mistake, the manager, lawyer or accountant is ultimately on the hook," Savare warned. "Fiduciaries have ethical and professional obligations to their clients. AI represents a massive shift in our ability to harness computational power to improve processes and increase efficiency, but it has to be used responsibly and intelligently. There must be human oversight over the entire process."

Impact on Attorney Growth

Savare opined that overreliance on AI can also hurt attorneys' growth. "I do not think this risk is overblown at all. AI displaces the type of work that a junior associate, analyst or accountant typically does," he stressed. "Most in-house attorneys have a fair amount of experience, but that is not always the case. Unless an attorney works at a very small firm, no partner will ask him or her to litigate a case immediately after graduating law school."

Law school, Savare continued, teaches a lot of theory but often few fundamental skills, which means that junior attorneys must learn on the job. First- and second-year litigation attorneys, for instance, typically spend a lot of time on document discovery, during which they must determine whether documents are responsive to a plaintiff's demands or whether any communications are privileged and should not be disclosed. "A lot of thinking goes into that, even though the tasks are routine. A junior attorney learns the process of litigating a case," he said.

"If you turn all of those routine tasks over to a computer, junior attorneys will not know how to do those tasks or why they are important," Savare explained. "They will never reach the critical mass of knowledge that is needed to reach the next level, which involves deeper analysis and strategizing. Not having those opportunities could harm their career growth, training and development."

Risk of Not Knowing Agreements

There is always a risk that lawyers will not remember enough about their own legal agreements, regardless of whether AI has been used or not. "Even an attorney who has been personally drafting and negotiating a contract for a client over the course of several weeks or months will forget a lot about that agreement; he or she will not be touching that contract every day," said Savare. "Of course, that is not an issue as long as the attorney has the time to adequately refresh his or her recollection about what the contract says."

Nevertheless, AI can exacerbate this issue. By way of analogy, Savare pointed to people's reliance on phones in the modern age. "If I asked you to give me the phone numbers for all of your friends in your address book, you likely would not be able to," he said. "The human mind only has so much space. Things that we do not need to remember – *i.e.*, things that we can offload to a computer – do not occupy our minds. Issues arise, however, when the computer is not around." He continued that, in his professional capacity, he would never want to be so reliant on technology that he would not be able to function without it. "You should always be able to reverse engineer the work product and do it the old-fashioned way."

Issues With Interpretation and Intent

Overreliance on AI can also lead to issues pertaining to contractual interpretation and intent. "I am unaware of any AI in the legal world that can show its intent or rationale," Savare said. "To the extent that it can be done, it is predicated on a human explaining to the computer through its programming – that is, through if-then statements."

For example, if a fund manager wants to enter into a perpetual software license as the licensee, it will want to have a source-code escrow provision that provides that if the licensor goes out of business, it will give the licensee the source code. A programmer can code the AI to add a source-code escrow provision if the manager is entering into a perpetual software license as the licensee and no such provision currently exists.

"In essence, the AI could explain its rationale for adding the provision by pointing to that if-then statement in the code. This is different than the AI divining the rationale on its own," explained Savare. "Although there are programs that can analyze data and make predictions – for instance, Lowenstein uses software that analyzes a judge's prior opinions, identifies his or her tendencies based on those opinions and makes probabilistic predictions based on the facts of the current case – that is not the same as explaining intent."

"These issues can be problematic for fund managers, which operate in a heavily regulated industry, because it is often very important for a manager to explain why it made certain decisions," cautioned Savare.

Problems with explainability may be exacerbated as AI becomes more sophisticated. For instance, AI will likely take greater control over the contracting process in the future, allowing managers to automate the construction of, among other things, private placement memoranda, limited partnership agreements and subscription agreements. "Eventually, AI will be able to leverage the manager's database of contracts – both executed and draft – to prepare and negotiate provisions that align with that manager's strategy, business plan, etc.," he said.

Recommendations for GCs on Integrating AI

GCs should begin by hiring the right people, Savare suggested – specifically, attorneys who are intelligent, capable and want to learn. Next, they should create programs – both formal and informal – to train those people. "An open-door policy between leaders and new hires is tremendously important; if a new hire has a question, he or she should not be intimidated to ask it," he said. "Conversely, senior personnel need to understand that they can always learn, too. It is very possible that junior attorneys know more about the newest technologies, and that is something they can teach."

"Too often, partners or other seasoned attorneys do not want to use AI or other technology because they are set in their own ways," Savare continued. "Thus, it is very important to generate buy-in from leadership across the organization." Leaders should cogently and persuasively explain why AI is important generally and why it should be used at the organization specifically, including how it will benefit the firm and its clients, as well as the expected return on investment due to increased efficiency, decreased expenses or other reasons.

"That also means creating a plan to do things correctly and ensuring that everyone is comfortable with the technology," Savare added. Ultimately, technology should not be integrated into a manager's infrastructure for the sake of integration; it should be used as a solution to best service clients, shareholders or investors. "People will always be more important – the machines are there to help them and their businesses grow."

Proper integration means that a manager should conduct pilot programs to ensure that the candidate software packages function as intended. According to Savare, larger funds evaluate multiple packages simultaneously to determine what works best for the organization. This allows the manager to be more selective, deliberative and comprehensive in its search and analysis.

"Managers need to apply high levels of due care and scrutiny when analyzing AI vendors and their software packages, because a seemingly small mistake can result in large losses," said Savare. "They must exhaustively analyze their options. There are very few processes where only one company offers a solution." Larger managers, he continued, will often downselect to a couple of viable options after conducting pilots of several others; use those two for a year; and then compare the results side-by-side to determine which is better. Only then will they sign long-term deals and invest substantial sums of money on products.

"A manager will probably want to have a core group of users test it out first," recommended Savare. "If the testing goes well, it can be expanded to a larger group, including those who may not have the same technical capabilities." He added that the software should be rolled out slowly and that individuals be trained on it and rewarded for using it. "Regardless, the process needs to be methodical and transparent." In addition, he noted that it is very important for leadership to associate themselves with information technology professionals early on.

Ethical Considerations: Privacy

The U.S. Department of the Treasury report to President Trump on nonbank financials, financial technology and innovation warns that as data becomes more pervasive, "financial and nonfinancial data may be increasingly shared without . . . understanding and informed consent. Moreover, the power of AI and machine learning tools may expand the universe of data that may be considered sensitive as such models can become highly proficient in identifying users individually."

For example, AI and machine learning models may leverage public data in certain circumstances, such as with anti-money laundering/know your customer compliance, which would, according to the Financial Stability Board report on AI and machine learning in financial services, make it "necessary to consider how the output . . . should be protected, while protecting the anonymity of each [data subject]."

"Using AI may increase the types and amount of data that managers collect. Thus, they must reassess whether they are complying with the privacy rules in the jurisdictions in which they operate, as well as in the jurisdictions where they are doing business," said Michelle Ann Gitlitz, founder and co-chair of Blank Rome's blockchain technology and digital currencies group, who also has experience advising hedge funds in other emerging technologies, including AI. "This may mean compliance with the E.U. General Data Protection Regulation (GDPR) or the California Consumer Privacy Act, which will become effective in January 2020."

Article 22 of the GDPR states that data subjects "have the right not to be subject to a decision based solely on automated processing . . . which produces legal effects concerning him or her or similarly affects him or her." In addition, Recital 71 of the GDPR states that a data subject shall have the right to "obtain an explanation of the decision reached" by an algorithm. It is unclear, however, whether the so-called "right to explanation" is binding given that it does not appear in the text of the regulation itself.

See our two-part series on the impact of, and compliance with, the GDPR: Part One (Feb. 21, 2019); and Part Two (Feb. 28, 2019).

"I am not sure that these privacy concerns can be solved by regulation alone," opined Michael Gallina, CEO and founder of Phylum Data Suite, a technology firm that offers AI products to the asset management industry and others. "It seems rather utopian and unrealistic to imagine that there will not be 'bad actors' or businesspeople inclined to take advantage of private data. Individuals for whom privacy is a major concern will, therefore, want to be responsible for protecting themselves. The government can only do so much to stifle such a complex and dynamic problem."

"The broader picture is that our technology systems as a whole are generating a massive amount of data exhaust across categories," reflected Lex Sokolin, global financial technology co-head at ConsenSys who has assisted hedge funds and others with their strategic decisions regarding, among other technologies, AI. "This includes media, health, commerce and financial data about every individual at scale. As a result, if an individual does not own his or her identity and the data attached to it, then machine learning systems are able to impute a meaningful amount of information without touching the regulated data itself." The Chinese technology giants, for example, have been able to build "financial services behemoths" by combining all these data assets together.

Sokolin agreed that regulation alone cannot solve the problem but argued that it can go several steps toward highlighting what the public is interested in protecting. "The GDPR and the Revised Payment Service Directive in Europe have shown the power of public institutions enforcing conversations upon the financial services industry. Tech companies have been very good at getting around regulation, however, by partnering with willing regulated participants, like banks and investment firms."

"When looking more specifically at the investment management industry and hedge funds, the question of who holds what kind of data and for how long is key," Sokolin added. "I am less worried about particular investment firms, and more worried about horizontal players like global custodians and data aggregators, which are accessing much more systemic data."

When asked how hedge funds can minimize these concerns, Sokolin stated that they should gain an understanding of how their service providers use the data generated by their trading and their customer information. "While there is unlikely to be something explicitly malicious, it is increasingly possible to impute information based on its surrounding factors," he continued. "Firms should develop an understanding of the data policies – regarding both retention and use – of their vendors."

"Creating a vertical that specializes on the ethics of AI is a good start for fund managers," added Gallina. "There needs to be more subject matter expertise in this space. The creation of a 'department' to address ethics can help spur that and give society more clarity on whether this is something that can – or should – be managed organizationally or via the government."

"In terms of protecting privacy, ensuring that data is properly anonymized is a simple first step," continued Gallina. "There are a number of well-established platforms that do this. On the other side, investors should have discussions with companies that specialize in data privacy, which can provide investors with sufficient background to know the right questions to ask funds."

Employee privacy is another story, however. "The assumption that there is no privacy in investment firms is fairly well established," remarked Sokolin. "Much of the relevant activity of financial professionals in a regulated firm has to be retained anyway, so I would expect powerful monitoring software and profiling tools. It still surprises me when there are attempts at insider trading using company tools – the guards against it are increasingly powerful."

"What may be intrusive, however, is the ability of these monitoring tools to pull on social media and other external sources – essentially coopting the private lives of employees," Sokolin warned. "There has not been a great solution to how financial professionals can and should use social media. From what I have seen, the answer is that all personal content becomes subject to the same retention requirements, and AI scans if the person opts into using those tools at work."

For more on social media, see "How to Navigate the Testimonial Rule in the Age of Social Media: Handling Clients' Online Reviews" (Aug. 2, 2018); and "How Can Fund Managers Address the Regulatory, Compliance, Privacy and Ethics Issues Raised by Social Media?" (Nov. 21, 2012).