

Machine Learning: A Revolution in Risk Management and Compliance?

Bart van Liebergen – Associate Policy Advisor, Institute of International Finance

Abstract

Machine learning and artificial intelligence are big topics in the financial services sector these days. Financial institutions (FIs) are looking to more powerful analytical approaches in order to manage and mine increasing amounts of regulatory reporting data and unstructured data, for purposes of compliance and risk management (applying machine learning as “RegTech”) or in order to compete effectively with other FIs and FinTechs. This article aims to give an introduction to the machine learning field and discusses several application cases within financial institutions, based on discussions with IIF members and technology ventures: credit risk modeling, detection of credit card fraud and money laundering, and surveillance of conduct breaches at FIs.

Two tentative conclusions emerge on the added value of applying machine learning in the financial services sector. First,

the ability of machine learning methods to analyze very large amounts of data, while offering a high granularity and depth of predictive analysis, can improve analytical capabilities across risk management and compliance areas in FIs. Examples are the detection of complex illicit transaction patterns on payment systems and more accurate credit risk modeling. Second, the application of machine learning approaches within the financial services sector is highly context-dependent. Ample, high-quality data for training or analysis are not always available in FIs. More importantly, the predictive power and granularity of analysis of several approaches can come at the cost of increased model complexity and a lack of explanatory insight. This is an issue particularly where analytics are applied in a regulatory context, and a supervisor or compliance team will want to audit and understand the applied model.

INTRODUCTION

In recent years, machine learning and artificial intelligence have seen increasing interest and popularity in the financial services community, as hopes are that they can dramatically improve analytical capabilities and streamline and automate all kinds of business lines including credit underwriting, compliance, interaction with clients, and risk management. The Institute of International Finance (IIF) has previously written about the use of machine learning/AI as “RegTech” in banking, and in the new business models of FinTech.¹

In past years, the amounts of data gathered in financial institutions (FIs) have increased significantly as the detail of reporting requirements has mushroomed and digitalization of services is creating a large amount of high-frequency, unstructured consumer data. As a result, FIs have a clear need for more powerful analytical tools to cope with large amounts of data of all kinds of sources and formats, while maintaining or improving granularity of analysis. Machine learning is widely seen in the financial services sector as a technique that may deliver that analytical power. It is a subfield of statistics that quickly gained prominence in the tech community in recent years. While elements of machine learning go back to the early 20th century, widespread use picked up as computing innovations and greater availability of high-frequency data allowed it to model complex, non-linear relationships, while making machine learning much easier to be applied.

This article aims to shed more light on the concept of machine learning and its uses within financial services: machine learning’s links with other types of statistical analysis, its possibilities, and its limits. It will also briefly touch on deep learning, a form artificial intelligence that has its roots in machine learning. Thereafter, applications within banking will be discussed through three use cases of machine learning: credit risk modeling, detection of fraud and money laundering, and surveillance of conduct breaches and abusive behavior within financial institutions.

BACKGROUND TO MACHINE LEARNING

Machine learning comprises a broad range of analytical tools, which can be categorized into “supervised” and “unsupervised” learning tools. Supervised machine learning involves building a statistical model for predicting or estimating an output based on one or more inputs (e.g., predicting GDP growth based on several variables). In unsupervised learning,

a dataset is analyzed without a dependent variable to estimate or predict. Rather, the data is analyzed to show patterns and structures in a dataset.²

Machine learning is a particularly powerful tool for prediction purposes. By identifying relationships or patterns in a data sample, it is able to create a model incorporating those relationships that lead to the most powerful out-of-sample predictions. Such a model is created by running variables and the model on subsamples of the data to identify the most powerful predictors, and then testing the model on many different data subsamples.³ This can be done thousands of times so that the model can “learn” from the data and improve its predictive performance. Due to its reliance on large datasets and heavy computing power, machine learning is closely associated with the “big data revolution.” In all, “[t]he extraordinary speed-up in computing in recent years, coupled with significant theoretical advances in machine-learning algorithms, have created a renaissance in computational modeling.”⁴

The accuracy of some supervised machine learning approaches is further augmented through their ability to conduct non-parametric analyses, which can flexibly fit any model to estimate the data. This is in contrast to some conventional statistical approaches that start out by making an assumption about the relationship between the dependent and independent variable. Linear regression, for example, assumes that this relationship is linear, while this does not necessarily need to be the case. Several machine learning approaches, in contrast, are also able to infer non-linear relationships, which makes them better able to fit the data.

1 See IIF, “Regtech in financial services: technology solutions for compliance and reporting,” March 2016, at <http://bit.ly/2mD0q39>. Also, IIF, “Digitizing intelligence: AI, Robots and the future of finance” at <http://bit.ly/1WIEbE6>.

2 James, G., D. Witten, T. Hastie, and R. Tibshirani, 2013, *An introduction to statistical learning: with applications in R*, Springer Texts in Statistics. The difference between both methods has also been described as the supervised ML being based on “labeled” data to train the algorithm, while unsupervised ML lacks training data with such labels and has to determine correlations by itself. However, this is the same as having a dependent variable or not: labels in the training data are values of the dependent variable.

3 Large datasets are typically divided into several separate samples to estimate a model (training), to choose the model (validation), and to evaluate how well the chosen model performs (testing).

4 Khandani, E. A. J. Kimz, and A. W. Lox, 2010, “Consumer credit risk models via machine-learning algorithms,” *Journal of Banking & Finance* 34, 2767–2787

Machine learning methods

The machine learning spectrum comprises many different analytical methods, whose applicability varies with the types of statistical problem one might want to address. Broadly speaking, machine learning can be applied to three classes of statistical problems: regression, classification, and clustering. Regression and classification problems both can be solved through supervised machine learning; clustering is an unsupervised machine learning approach.

Regression problems involve prediction of a quantitative, continuous dependent variable, such as GDP growth or inflation. Linear learning methods try to solve regression problems including partial least squares⁵ and principal component analysis; non-linear learning methods include penalized regression approaches, such as LASSO and elastic nets.⁶ In penalized approaches, a factor is typically added to penalize complexity in the model, which should improve its predictive performance.

Classification problems typically involve prediction of a qualitative (discrete) dependent variable, which takes on values in a class, such as blood type (A/B/AB/O). An example is filtering spam e-mail, where the dependent variable can take on the values SPAM/NO SPAM. Such problems can be solved by a decision tree, “which aims to deliver a structured set of yes/no questions that can quickly sort through a wide set of features, and thus produce an accurate prediction of a particular outcome.”⁷ Support vector machines also classify observations, but by applying and optimizing a margin that separates the different classes more efficiently.⁸

In clustering, lastly, only input variables are observed while a corresponding dependent variable is lacking. An example is exploring data to detect fraud without knowing which observations are fraudulent and which not. An anti-money laundering (AML) analysis may nonetheless yield insights from the data by grouping them in clusters according to their observed characteristics. This may allow an analyst to understand which transactions are similar to others. In some instances, unsupervised learning is first applied to explore a dataset; the outputs of this approach are then used as inputs for supervised learning methods.⁹

Table 1 classifies popular machine learning approaches according to their (un)supervised learning character, and the types of problems they can be applied to.

Prediction versus explanation

Machine learning’s ability to make out-of-sample predictions does not necessarily make it appropriate for explanation or

	Linear methods	Non-linear methods	
Problem type	Supervised		
	Regression	<ul style="list-style-type: none"> • Principal components • Ridge • Partial least squares • LASSO 	Penalized regression: <ul style="list-style-type: none"> • LASSO • LARS • elastic nets Neural networks and deep learning
	Classification	Support vector machines	Decision trees: <ul style="list-style-type: none"> • classification trees • regression trees • random forests Support vector machines Deep learning
	Unsupervised		
	Clustering*	Clustering methods: K- and X-means, hierarchical Principal components analysis Deep learning * Since unsupervised methods do not describe a relation between a dependent and interdependent variable, they cannot be labelled linear or non-linear.	

Table 1 – Overview of machine learning methods

inference as well, as statistical methods are typically subject to a trade-off between explanatory and predictive performance. A good predictive model can be very complex, and may thus be very hard to interpret.¹⁰ For predictive purposes, a model would need only to give insight in correlations between variables, not in causality. In the case of credit scoring a loan portfolio, a good inferential model would explain why certain borrowers do not repay their loans. Its inferential performance can be assessed through its statistical significance and its goodness-of-fit within the data sample. A good predictive model, on the other hand, will select those indicators that prove to be the strongest predictors of a borrower default. To that end, it

5 PLS is used to find the fundamental relations between two matrices through linear regression.

6 LASSO stands for least absolute shrinkage and selection operator. LASSO and elastic nets both perform variable selection, yet apply different types of penalties for model complexity.

7 Tiffin, A., 2016, “Seeing in the dark: a machine-learning approach to nowcasting in Lebanon,” IMF Working Paper WP/16/56, March

8 Auria, L., and R. Moro, 2008, “Support vector machines as a technique for solvency analysis,” Deutsches Institut für Wirtschaftsforschung, discussion papers no. 811, Berlin.

9 Bholat, D., S. Hansen, P. Santos, and C. Schonhardt-Bailey, 2015, “Text mining for central banks: handbook,” CCBS Handbook No. 33, Bank of England

10 Tiffin (2016)

does not matter whether an indicator reflects a causal factor of the borrower's ability to repay, or a symptom of it. What matters is that it contains information about the ability to repay.

Tackling overfitting: bagging and ensembles

Excessively complex models can also lead to "overfitting," where they describe random error or noise instead of underlying relationships in the dataset. Model complexity can be due to having too many parameters relative to the number of observations.¹¹ In machine learning, overfitting is particularly prevalent in non-parametric, non-linear models, which are also complex by design (and therefore also typically hard to interpret). When a model describes noise in a dataset, it will fit that one data sample very well, but will perform poorly when tested out-of-sample.¹²

There are several ways to deal with overfitting and improve the forecast power of machine learning models, including "bootstrapping," "boosting" and "bootstrap aggregation" (also called bagging).¹³ Boosting concerns the overweighting of scarcer observations in a training dataset to ensure the model will train more intensively on them. For example, one may want to overweight the fraudulent observations due to their relative scarcity when training a model to detect fraudulent transactions in a dataset. In "bagging," a model is run hundreds or thousands of times, each on a different subsample of the dataset, to improve its predictive performance. The final model is then an average of each of the run models. Since this average model has been tested on a lot of different data samples, it should be more resilient to changes in the underlying data. A "random forest" is an example of a model consisting of many different decision tree-based models.

Econometricians can take this concept even further by combining the resulting model with a model based on another machine learning technique. The result is a so-called ensemble: a model consisting of a group of models whose outcomes are combined by weighted averaging or voting.¹⁴ It has been shown that averaging over many small models tends to give better out-of-sample prediction than choosing a single model.¹⁵

A theory-free approach to analysis?

Due to a typical lack of explanatory power and inherent complexity of machine learning models, the discipline has been criticized by some as "a theory-free analysis of mere correlations," which is "inevitably fragile."¹⁶ Machine learning relies on found in-sample (past) correlations to predict out-of-sample (future) correlations, without always offering an understanding of the relationship analyzed. In that sense, it is as much a backward-looking way of prediction as other statistical

approaches. It can only be more accurate at inferring those correlations. However, one observer has noted, "[i]f you have no idea what is behind a correlation, you have no idea what might cause that correlation to break down."¹⁷

Deep learning and neural networks: from machine learning to artificial intelligence

So far, discussion has focused on "classic" machine learning methods that are applied to statistical problems with well-defined and structured datasets. Additionally, machine learning approaches have been advanced and combined to solve all kinds of complex problems, functioning as "artificial intelligence." One of the dominant approaches is deep learning, a learning approach that can be based on both supervised and non-supervised methods; all are non-linear.¹⁸

In deep learning, multiple layers of algorithms are stacked to mimic neurons in the layered learning process of the human brain. Each of the algorithms is equipped to lift a certain feature from the data. This so-called representation or abstraction is then fed to the following algorithm, which again lifts out another aspect of the data.¹⁹ The stacking of representation-learning algorithms allows deep-learning approaches to be fed with all kinds of data, including low-quality, unstructured data; the ability of the algorithms to create relevant abstractions of the data allows the system as a whole to perform a relevant analysis. Crucially, these layers of features are not designed by human engineers, but learned from the data using a general-purpose learning procedure.²⁰

11 For example, R^2 , a goodness-of-fit indicator, tends to increase (and cannot decrease) with any variable that is added to the model, whether or not it makes sense in the context. See Ramanathan, R., 2002, "Introductory econometrics with applications, South-Western

12 James et al. (2013), p. 22.

13 In regression models, overfitting is also mitigated through "ridge regression" or "LASSO," both of which add a factor penalizing complexity from having too many variables.

14 Tiffin (2016).

15 Varian, H., 2014, "Big data: new tricks for econometrics," *Journal of Economic Perspectives* 28:2, 3-27

16 Harford, T., 2014, "Big data: are we making a big mistake?" *Financial Times*, March 28

17 Ibid.

18 Najafabadi, M. M., F. Villanustre, T. M. Khoshgoftaar, N. Seliya, R. Wald, and E. Muharemagic, 2015, "Deep learning applications and challenges in big data analytics," *Journal of Big Data* 2:1

19 LeCun, Y., Y. Bengio, and G. Hinton, 2015, "Deep learning," *Nature* 521:28, 436-444

20 Ibid.

Deep learning is being applied to a wide range of uses. The ability to crunch large amounts of raw data and to identify complex patterns in it makes it particularly well-placed to analyze “big data,” such as the user datasets of tech giants, such as Google, Microsoft, and Amazon.

Given that it was partly developed by the U.S. National Security Agency, it is perhaps unsurprising that deep learning has proved to be very proficient at face recognition and natural language understanding, including question answering and language translation. Upon “overhearing” a discussion, it is able to classify the topic of discussion and the sentiments of the speakers.²¹ While some conventional machine-learning approaches can be equipped to solve non-numeric problems as well (for example, x-means clustering has been applied to text mining), deep learning has often proved to be more accurate. However, a typical deep-learning system is extremely complex and requires a dataset with hundreds of millions of labeled observations only to be trained. In many fields, availability of sufficient data for such extremely large datasets is hardly a given.

Application within financial services

In past years, the amounts of data gathered in financial institutions (FIs) have increased significantly as the details of reporting requirements have mushroomed and digitalization of services is creating a large amount of high-frequency, unstructured consumer data. As a result, FIs have a clear need for more powerful analytical tools to cope with large amounts of data of all kinds of sources and formats, while maintaining or improving granularity of analysis.

After the financial crisis of 2008-09, many new regulations and supervisory measures were introduced that required FIs to report more detailed and more frequent data on more aspects of their business models and balance sheets. Under the new capital regime, banks report large exposures, liquidity measures, collateral, and capital levels. Stress tests are based on all kinds of firm data including loan-level balance sheet data and qualitative aspects of the business model. The Federal Reserve’s CCAR exercise requires FIs to consider the impact of more than 2000 economic variables on their business. For insurers, Solvency II has dramatically increased reporting requirements.

These processes create large amounts of reporting data that need to be well-defined and structured, aggregated across the group, and delivered in-time with supervisors. Regulators have, therefore, introduced numerous initiatives to improve the quality of supervisory data and the ability of financial

institutions to deliver these data. The Basel Committee’s Principles for Risk Data Aggregation (Basel 239) sets standards for G-SIBS to improve their IT systems and reporting structures. IFRS 9 aims to improve the quality of supervisory data.

Apart from reporting data, FIs are increasingly able to gather large amounts of low-quality, unstructured, high-frequency data. These include outputs from consumer apps and other digital interactions with clients, metadata from payment systems, and external data sources, such as social media feeds, which can be mined to gauge insights on market sentiment. This type of data is typically called “big data.”

With practically all aspects of FI’s business model regulated and supervised with detailed risk metrics, running a bank, insurer, or asset manager is increasingly becoming a matter of optimization within hundreds of constraints. To compete effectively, they need to find this optimum while also mining consumer data for detailed insights on client preferences and behavior.

The extensive set of machine learning approaches is well situated to deliver this analytical power in different contexts due to its ability to cope with (or better said, its need for) extremely large datasets and the granularity of analysis. For the mining of high-quality, structured supervisory data, more conventional machine learning techniques are typically applied. To mine high frequency, low quality “big data” sources, Google-like deep learning and neural network techniques are applied, which cope with these data due to their representation learning abilities.

Below, the state of play in three use cases of machine learning is being discussed: the modeling of credit risk, detection of fraud and money laundering, and the detection of conduct risk and abusive behavior within financial institutions.

THREE USE CASES

Credit risk and revenue modeling

Since the early 2000s, an extensive academic literature on the use of machine learning methods to model credit risk has developed. To give just a few examples, Angelini et al. (2007) apply a neural network approach to model SME credit risk on

²¹ Ibid.

a small dataset of Italian SMEs. Auria and Moro (2008) assess company solvency using support vector machines, and find that they produce more accurate out-of-sample predictions than existing techniques. Khandani et al. (2010) apply generalized classification and regression trees (CART) to a large dataset of a commercial bank to build consumer credit risk models. These combine traditional credit factors, such as debt-to-income ratios, with consumer banking transactions, which greatly increases the predictive power of the model.

FIs have traditionally used linear, logit, and probit regressions to model credit risk for capital requirements, stress-testing, and internal risk management procedures.²² Recently, many have started to experiment with the application of machine learning methods to improve financial risk predictions. Unsupervised methods are typically used to explore the data, while regression and classification methods (trees, support vector machines) can predict key credit risk variables as probability of default or loss-given default. Banks normally have extensive records of loan-level data to serve as inputs.

Banks have sometimes also experienced that machine learning can be hard to apply, as methods can be complex and models sensitive to overfitting the data. Thereby, the quality of data within banks is not always fit enough for advanced statistical analysis, while banks are not always able to consolidate the data from across the financial group, among others, due to inconsistent data definitions across jurisdictions and the use of multiple systems. Non-parametric and non-linear approaches (support vector machines, neural networks, and deep learning) and ensembles are so complex that they are practically “black boxes” that are hard, if not impossible, for any human to understand and audit from the outside. That makes these models hardly useful for regulatory purposes, such as the development of internal models in the Basel Internal Ratings-Based approach. Financial supervisors typically require risk models to be clear and simple in order to be understandable and verifiable and appropriate for validation by them.

That does not, however, rule out the use of machine learning to optimize parameters and models with a regulatory function. Linear and simple non-linear machine learning approaches can be applied and still perform better than similar non-machine learning approaches. Machine learning can also be applied to select variables and optimize parameters in existing, linear regulatory models. Khandani et al (2010) stress that CART (tree) models produce easily interpretable decision rules whose logic is clearly laid out, despite their non-linear character. Indeed, there have been examples already of banks applying machine learning in a regulatory context. In a public

example, Citigroup hired an external vendor to build a revenue forecasting model for the 2015 CCAR exercise.²³

Fraud

One area in which machine learning has been applied for more than a decade and with significant success is the detection of credit card fraud. Banks have equipped their credit card payments infrastructures with monitoring systems (so-called workflow engines), which monitor payments for potential fraudulent activity. Fraudulent transactions can then be blocked in real-time. The fraud models used by these engines have been trained on historical payments data.

The high frequency of credit card transactions provides the large datasets required for algorithm training, back testing and validation. Furthermore, since banks are able to verify unambiguously which transactions were fraudulent and which were not, they can construct clear historical data with relevant fraud and non-fraud labels to train classification algorithms. The historical transaction datasets showcase a wide variety of pre-determined features of fraud, which distinguish normal card usage from fraudulent card usage, ranging from features from transactions, the card holder, or from transaction history.

The detection of money laundering and terrorism financing through payments systems stands as a contrast to machine learning’s long-standing record in credit card fraud. Many banks are still relying on conventional rules-based systems, which focus on individual transactions or simple transaction patterns. These systems are often unable to detect complex patterns of transactions or obtain a holistic view of transactions behavior on payment infrastructures. Due to their coarse selection methods, the number of false positives created by these systems is substantial. As a result, significant human capacity is required for the assessment of alerts and filtering false positives from actual suspicious observations. In addition, impediments to data sharing and data usage, as well as long-established regulatory requirements, have complicated innovation in the AML/CFT area.²⁴

22 In a probit model, the dependent variable is binary (can only take two values); in a logit model, the dependent variable is categorical.

23 Ayasdi, “CCAR stress test,” <http://bit.ly/2m5n4y2>, undated; and “After yesterday, CCAR less stressful for Citigroup,” March 6, 2015, <http://bit.ly/2mmbfph>

24 See the IIF’s forthcoming report on the use of “regtech for AML” and submissions to FATF and the BCBS for more information on data sharing issues in AML/CFT on www.iif.com.

Machine-learning systems have the potential to improve detection of money laundering activity significantly, due to their ability to identify complex patterns in the data and combine transactions information at network speed, with data from many other sources to obtain a holistic picture of a client's activity. Indeed, these systems have already been shown to bring false positives down significantly.²⁵

However, application so far in the AML space has lagged for several reasons. First, money laundering is hard to define. There is no universally agreed definition of money laundering and financial institutions do not receive feedback from law enforcement agencies on which of their reported suspicious activities have turned out to be money laundering. It is, therefore, more difficult to train ML-detection algorithms using historical data, because an incidence of money laundering typically is not firmly established. As a second-best, FIs are optimizing ML detection algorithms using lower-level suspicious activity reports as a depending variable for classification – using classification between alerts that the bank could classify as false alerts, and those that moved on to be submitted as SARs to law enforcement agencies.

Unsupervised learning methods are also applied to AML/CFT as they “learn” relevant patterns from the data by clustering transactions or client activity. This yields additional insights, since laundering methods take all kinds of form and develop on a continuous basis.

An example of such unsupervised learning is clustering. Clustering requires large datasets where it can automatically find patterns within the data without the need for labels. Clustering works by identifying outliers as points without any strong membership in any one cluster group, thus finding anomalies within subsets of the data. In AML, clustering is one of the methods used to group together data; using other analytics, such as topological data analytics and dimensionality reduction, machine learning can reduce the significant amounts of false positives often associated with alternative methods.

Surveillance of conduct and market abuse in trading

A third area in which machine learning is increasingly being applied within financial institutions is the surveillance of conduct breaches by traders working for the institution. Examples of such breaches include rogue trading, benchmark rigging, and insider trading – trading violations that can lead to significant financial and reputational costs for FIs. In the last couple of years, automated systems have been developed that monitor the behavior of traders in multiple ways and with increasing accuracy.

The capabilities of the first generation of these surveillance systems were limited to monitoring trading behavior, and only through assessing single trades. However, the improved ability of machine learning approaches to identify large, complex patterns in data has allowed a new generation of systems to analyze entire trading portfolios. These systems are also able to link trading information to other behavioral information of a trader, such as e-mail traffic, calendar items, building check in and check out-times, and even phone calls. Technologies, such as natural language processing (typically based on deep learning) and text mining (which can be based on several learning algorithms²⁶), have made those sources machine-readable and suitable for automated analysis. The outputs of the trading behavior and communications of one or multiple traders are then integrated and compared to a profile of “normal” behavior. When a trader's behavior or trading performance deviates from what is deemed normal, the system will send an alert to the FI's compliance team.

There are several challenges to applying machine learning in this space. First, there are typically no labeled data to train algorithms on, as it is legally complex for financial institutions to share the sensitive information on past breaches with developers. Supervisory learning approaches are, therefore, hard to apply. Second, a surveillance system needs to be auditable for supervisors and for compliance officers, and needs to be able to explain to a compliance officer why certain behavior has set off an alert. For systems that are entirely based on machine learning, that can be difficult due to the “black box” character of learning approaches. In order for an alert to be interpretable and actionable for compliance teams, it should ideally be linked to detection of a specific kind of behavior, rather than based solely on a statistical correlation in the data.

These issues can be addressed at least partly by founding the learning system on a behavioral science-based model, which incorporates human decisions and behavioral traits. In a way, such a model addresses the lack of explanatory power of machine learning approaches. Any alerts from the system will be based on deviations it has identified from the model. However, the inclusion of machine learning approaches on top of the model creates a feedback loop in the system through which it can adapt to evolving behavior, and “get to know” a

25 Adamson, D., 2016, at “Machine learning – the future of compliance?” panel discussion at Sibos conference, September 28

26 Bholat et al., 2015.

trader as it ingests more data. That is a crucial difference with previous rules-based systems, which are unable to tailor their surveillance methods to changed probability distributions and correlations. Consequently, these systems are typically based on more conventional types of machine learning, which can be audited and explained more easily than complex types, such as neural nets and deep learning.

A practical barrier to the implementation of automated surveillance systems is the fragmentation and complexity sometimes found in FI's IT systems. To gain a perspective on a trader's behavior, surveillance systems require information from many sources, which are likely to be found in different systems that can be mutually incompatible or slow to deliver.

CONCLUSION

Machine learning and artificial intelligence are big topics in many fields of business these days, including the financial services sector. FIs are looking to more powerful analytical approaches as they need to manage and mine increasing amounts of regulatory reporting data and unstructured data, either for compliance purposes or in order to compete effectively with other FIs and FinTech's. There seems to be no aspect of the FI business model that is not impacted in some way by machine learning and artificial intelligence: it could improve insights into client preferences, risk management, the detection of fraud, and conduct breaches, and automate client support or allow for automated identity verification when coupled with biometrics.

This article has given an introduction to the machine learning field and has discussed several cases of application within financial institutions, based on discussions with IIF members and technology vendors: credit risk modeling, detection of credit card fraud and money laundering, and surveillance of conduct breaches at FIs. Two tentative conclusions emerge on the use of machine learning in the financial sector – tentative, because the field is developing fast and many FIs are still experimenting with machine learning in some spaces.

First, machine learning comprises a range of statistical learning tools that are generally able to analyze very large amounts of data while offering a high granularity and depth of analysis, mostly for predictive purposes. The ability of some approaches to infer non-linear relationships and to conduct data analysis without making assumptions about the shape or form of the relationship between variables (i.e., non-parametric) increases

the detail with which data can be analyzed and outcomes predicted. Unsupervised approaches allow for exploration of data without a dependent variable. Running algorithms thousands of times on training data and combining models improves their predictive power while limiting overfitting and maintaining analytical granularity.

Such improved, often automated, analytical capabilities allow FIs to gain better insights in business processes such as lending, risk management, customer interaction, and payments. With ever more data produced in these processes, machine learning can discover richer, more complex patterns and relationships as in the analysis of transactions or credit risk, or by connecting different datasets to draw more accurate overarching conclusions, as in the monitoring of conduct breaches.

Second, the application of machine learning approaches within the financial sector is highly context-dependent. Ample, high-quality data for training or analysis are not always available in FIs. More importantly, the predictive power and granularity of analysis of several approaches can come at the cost of increased model complexity and a lack of explanatory insight. This is an issue particularly where analytics are applied in a regulatory context, and a supervisor or compliance team will want to audit and understand the applied model. Fortunately, simpler machine learning approaches do exist, combining non-linear analysis with simplicity. Indeed, vendors of machine learning analytics in finance typically aim to combine machine learning's depth of insight with model simplicity, or add factor models to improve the auditability of their products. As it seems, there is an algorithm for every problem.