

# AI and machine learning for risk management

## Authors:

Saqib Aziz

*Rennes School of Business, France*

Michael Dowling

*Rennes School of Business, France*

## Abstract

We explore how artificial intelligence (AI) and machine learning solutions are transforming risk management. A non-technical overview is first given of the main AI and machine learning techniques of benefit to risk management. Then an analysis, using current practice and empirical evidence, is carried out of the application of these techniques to the risk management fields of credit risk, market risk, operational risk, and compliance ('RegTech'). We conclude with some thoughts on current limitations and views on how the field is likely to develop in the short- to medium-term. Overall, we present an optimistic picture of the role of AI and machine learning in risk management, but note some practical limitations around suitable data management policies, transparency, and lack of necessary skillsets within firms.

**Keywords:** AI, machine learning, risk management, RegTech, credit risk, operational risk, market risk

## 1 Introduction

Artificial intelligence (AI), and the machine learning techniques that form the core of AI, are transforming, and will revolutionize, how we approach financial risk management. Everything to do with understanding and controlling risk is up for grabs through the growth of AI-driven solutions: from deciding how much a bank should lend to a customer, to providing warning signals to financial market traders about position risk, to detecting customer and insider fraud, and improving compliance and reducing model risk. In this chapter we detail current AI and machine learning techniques being used and current applications of those techniques. We further envisage the future role for fully AI solutions as the natural next step after the widespread adoption of machine learning in helping the organization to manage risk.

An example of ZestFinance serves to illustrate the potential for AI and machine learning in risk management. ZestFinance was founded by a former Chief Information Officer of Google and in 2016 partnered with Baidu, the dominant search engine in China, to improve Baidu's lending decisions in the Chinese market. Baidu was particularly interested in making small loan offers to retail customers buying products from their platform. Unlike most developed countries, the risk with lending in the Chinese market is that less than 20 percent of people have credit profiles or credit ratings. Lending to people who have either 'thin' credit profiles, or no credit profiles, is inherently risky as there is no history to draw on to check the reliability of borrower. ZestFinance (with permission) taps into the huge volume of information on members held by Baidu such as their search or purchase histories to help Baidu decide whether to lend. They use thousands of data points per customer and are still able to make lending decisions on new applications in seconds. A reported trial in 2017 of their system led to a 150 percent increase in total small-item lending by Baidu with no increase in credit losses in the space of just two months<sup>1</sup>.

The exact nature of how ZestFinance makes these decisions is not disclosed except under the broad umbrella of machine learning and AI, but essentially what they will be using as a base is a core machine learning set of techniques around clustering and decision trees, and possibly deep learning. Thus a search history indicating accessing gambling websites, perhaps the use of virtual private networks, or even types of search terms used, would cluster a potential borrower into a higher risk group. While on the other hand, a history of online spending that supports the

---

<sup>1</sup> <https://www.technologyreview.com/s/603604/an-ai-fueled-credit-formula-might-help-you-get-a-loan/>

applicant's reported income in their application to borrow, or maybe even searches indicating upward career mobility, might cluster someone into a group of lower risk borrowers. The actual clusters will be significantly more refined than the simple examples given above. This approach crosses the line from machine learning to AI due to the automated nature of the lending decision making process.

The ZestFinance example illustrates the essence of how AI and machine learning can improve risk management. A standard credit score is largely a linear calculation of a small number (about 50) positive or negative numerical characteristics about a person and thus misses out on a huge amount of additional personal information that can help to either reduce negative risk or accept positive risk. This new information is often atypical and non-numerical; the type of data which machine learning is innately suited to analysing. Our chapter thus expands on this idea of how AI and machine learning can improve risk management – particularly around the techniques being used to make decisions based on such large volumes of atypical data. The next section briefly details, in a non-technical manner, the core machine learning techniques which can be applied to improve risk management. The remainder of the chapter is devoted to the actual application of AI and machine learning to various forms of risk management, finishing with a forward-looking perspective on what is next for the role of AI in risk management and some challenges that need to be addressed.

## **2 AI and machine learning techniques for risk management**

A first step is defining what we mean by AI and machine learning, and this is not necessarily a straightforward distinction. In a glib sense the public relations and fundraising functions of startups tend to use the more attractive AI term when they most often mean machine learning, but even in research there is a reasonably fluid distinction. One acceptable distinction is that machine learning is a core technique of AI involving learning from data, but that AI involves additional requirements. A full AI solution should be automated in terms of data identification, data testing, and making decisions based on the data testing. AI might also involve additional techniques in addition to machine learning, such as including hard-coded and logic rules. Machine learning on the other hand normally involves manual data identification and testing by the data scientist, and human decisions as to how to apply the outputted information. We can alternatively define AI in a less

technical manner, by saying that AI is simply intelligence demonstrated by machines, with intelligence being defined with reference to what we view intelligence as in humans. As it matters to finance we are normally particularly interested in artificial super-intelligence; that is machines that can demonstrate a finance-specific intelligence higher than human intelligence. Although a problem with this more general definition is that it doesn't clearly distinguish between AI and machine learning. Given the lack of technological and organizational readiness for pure AI and the reality that most claimed AI is in fact machine learning, in this section we outline the core machine learning techniques applied to risk management. In the following section and especially the last section we move our discussion more towards AI as the logical next step to follow from the widespread usage of machine learning techniques.

Machine learning falls into two broad categories of supervised and unsupervised machine learning<sup>2</sup>. In supervised learning you have input data that you wish to test to determine an output. This is similar to how in traditional statistics terms you have a range of independent variables that you test to determine relationship with the dependent variable. In unsupervised learning you only have input data and wish to learn more about the structure of the data. Figure 1 shows the distinction between these two categories as well as the main techniques within each category.

Regression machine learning is the closest group of techniques to that usually applied in traditional determination of the causal relationship between variables. In simple terms, we might describe a traditional linear regression equation for a credit lending risk assessment as perhaps the dependent variable being the risk of loan non-repayment, which is then sought to be explained by a range of independent variables that should influence the risk of loan non-repayment. These independent variables might, for example, include financial measures such average non-repayment rates, whether the person is full-time employed, whether they have a good credit history, and whether they own property.

Regression machine learning differs from traditional regression in that it uses regression techniques that allow for large numbers of variables to be used as independent variables and then

---

<sup>2</sup> As AI and machine learning are not a focus of this book, this section has been kept purposefully non-technical. An interested reader wishing to delve further into the technicalities is suggested at a first instant to read the accessible *The Master Algorithm* by Pedro Domingos, and for a thorough academic treatment of each of the methods described in this section to read *The Elements of Statistical Learning* by Trevor Hastie and colleagues. For practical learning the *Applied Data Science with Python* series of courses developed by the University of Michigan and *Machine Learning* course by Stanford University, all available through Coursera.org, are excellent resources.

automatically discarded if they lack explanatory power. This is a necessary feature due to the large range of data that is available to the data scientist. It also reduces the extent of theorizing needed to determine suitable independent variables. Thus LASSO regression zero weights independent variables with low explanatory power, while Ridge regression gives lower weights to variables in a model that are highly correlated with other variables in a model. In both cases the outcome is a reduced model that allows the data scientist to move from large numbers of potential explanatory variables to a smaller subset. A LARS regression works in the opposite direction to a LASSO and Ridge, by initially zero-weighting all variables and only adding variables that are shown to have explanatory power. Hu et al (2015) provide an interesting application of this set of techniques by applying a LASSO regression technique, among others, to the development of an automated investment trading advice system based on stock trend analysis and showing this method has improved trading performance compared to traditional methods.

Principal Component Analysis (PCA) and partial-least squares regressions are quite similar in that they both aim to reduce the number of variables by combining variables and extracting common factors. PCA is the more popular of the two as it is widely used in traditional statistics and therefore better understood. A simple example of PCA is that for a set of potential variables to be used to determine credit repayment risk consisting of: (1) owns a house, (2) owns a car, and (3) has significant savings: a common factor might be extracted from these that could be termed 'asset ownership'. A more applied example is provided by Zhong and Enke (2017) who use PCA to reduce 60 correlated economic and financial measures to a smaller set of factors and then apply a set of neural network machine learning techniques to forecast the US S&P 500 index. This latter example is a common use of PCA in machine learning: that it is used as a first step for dimension reduction and then a more advanced machine learning technique is applied to learn from the PCA components.

The other main category of supervised learning is 'classification', with support vector machines and decision trees being the most popular techniques within this group. The outcomes and visuals of decision trees are easily understood in practice and therefore amenable to explanation to non-data scientists. A famous application of the decision tree technique is to Titanic survival prospects. Starting with the initial tree trunk statistic that 62 percent of passengers died and 38 percent survived the Titanic disaster, we can build a decision tree to classify groups that had a greater or

lesser chance of survival. Thus a first classification of whether a passenger is male or female shows that only 19 percent of males survived while 74 percent of females survived. We can branch out the tree even further by, for example, looking at age groups beneath the first male-female division, and this shows that despite only 19 percent of males surviving, male boys under the age of 6 had a 67 percent survival chance.

This classification approach of creating sub-groups then helps to understand what characteristics determine outcomes. In practice decision trees for machine learning have many more challenges than in the Titanic example, particularly around the issue of overfitting to existing data and thus determined sub-groups having poor predictive power with new data and in new situations. Support vector machines (SVM) are more complex in formulation than decision trees but have the same essential end goal of creating groups based on input characteristics to classify and predict outcomes. In the case of SVM the approach is to map characteristics on a plane and classify groups based on similarity of where they are on this plane. As a financial risk example, Nazemi et al (2018) apply a SVM approach to predict financial loss to holders of bonds that have defaulted (normally some proportion of loss is recoverable from firm assets after bond default) and are able to demonstrate the superiority of the approach compared to more conventional methods.

Turning now to unsupervised learning<sup>3</sup>. We first have clustering machine learning techniques which have some similarity with SVM in that they involve mapping characteristics on a plane. The technique differs in that it is not trying to predict outcomes but instead create similar groups. Email ‘spam’ detection, for example, is usually based on a clustering approach – if an email looks like other emails that are deemed spam then it is likely to also be spam. K-means clustering is the most popular approach, although other methods such as X-means and hierarchical clustering are growing in popularity. We’ll focus on K-means clustering to describe the general technique. In K-means a number of cluster groups the data scientist wishes to arrive at is pre-determined (although in practice a range of the number of groups is tested), characteristics are mapped on a plane and a dividing line (not necessarily straight) is drawn that best distinguishes between groups. The idea of the iterative process behind the technique is to maximize the difference in means between determined groups.

---

<sup>3</sup> Deep learning will be discussed in this category, even though it can be either supervised or unsupervised

Deep learning and neural networks<sup>4</sup> are viewed as being at the forefront of machine learning techniques and are often classified separately to the machine learning techniques already described. The intuition behind these techniques is to more accurately model complex relationships between variables and ultimately to better mimic human decision making. To that extent these techniques represent the closest to actual AI techniques, albeit still missing some of the data identification and automation features necessary for true AI. Heaton et al (2017) provide a comprehensive analysis of the recent application of deep learning to finance, as well as highlighting the potential of the approach.

A key feature of deep learning is the addition of ‘hidden layers’ after the input data stage that allow multiple and combined influence between input variables to be determined by the modelling. As the input data progresses through the hidden layers variables are combined and recombined into newer factors weighted on influence from the prior layer. The addition of hidden layers between input and output is what creates the perceived problem with deep learning – that the process is a ‘black box’ – in that it is not always clear how inputs led to an output. This has obvious implications for use in risk management, where the very presence of a black box at the center of decision making can be its own source of risk in a firm.

### **3 AI and machine learning applications for risk management**

In this section we provide details and analysis of actual applications of AI and machine learning to various areas of risk management. We use a usual financial risk management distinction of dividing risk management areas into: credit risk, market risk, operational risk, and add a fourth category around the issue of compliance.

#### **3.1 Application to credit risk**

Credit risk is economic loss that emanates from the failure of a counterparty to fulfill its contractual obligations (e.g., timely payment of interest or principal), or from the increased risk of default

---

<sup>4</sup> While subject to some debate, it is worth thinking of deep learning as a newer term of neural networks, and the terms can in practice be used interchangeably

during the term of the transaction. Traditionally, financial firms have employed classical linear, logit, and probit regressions to model credit risk (Altman, 1968). However, there is now an increased interest by institutions in using AI and machine learning techniques to enhance credit risk management practices, partially due to evidence of incompleteness in traditional techniques. The evidence is that credit risk management capabilities can be significantly improved through leveraging AI and machine learning techniques due to its ability of semantic understanding of unstructured data.

The use of AI and machine learning techniques to model credit risk is not a new phenomenon though it is a growing one. Back in 1994, Altman and colleagues performed a first comparative analysis between traditional statistical methods of distress and bankruptcy prediction and an alternative neural network algorithm, and concluded that a combined approach of the two improved accuracy significantly (Altman et al, 1994).

It is particularly the increased complexity of assessing credit risk that has opened the door to machine learning. This is evident in the growing credit default swap (CDS) market where there is a lot of uncertain elements involving determining both the likelihood of an event of default (credit event) and estimating the cost of default in case a default takes place. Son et al (2016) use daily CDS of different maturities and different rating groups from January 2001 to February 2014 to show that nonparametric machine learning models involving deep learning outperform traditional benchmark models in terms of prediction accuracy as well as in proposing practical hedging measures.

The areas of consumer lending and SME lending involve large amounts of potential data and are increasingly relying on machine learning to make better lending decisions. The example of ZestFinance in the opening section is an example of this, but there are a wide number of similar firms operating in this space. There's a lot of empirical support for the effectiveness of machine learning. In consumer lending, Khandani et al (2010) propose a machine-learning technique based on decision trees and SVM that, when tested on actual lending data lead to cost savings of up to 25 percent. More recently, Figini et al (2017) show that a multivariate outlier detection machine learning technique improves credit risk estimation for SME lending using data from UniCredit Bank.



### 3.2 *Application to market risk*

Market risk is the risk that emanates from investing, trading, and generally from having exposure to financial markets. Kumar (2018) provides a structural overview of how machine learning can help in market risk management, noting the benefits at each stage from data preparation, to modelling, stress testing, and providing a validation trail for model explanation. A 2017 Financial Stability Board report '*Artificial Intelligence and Machine Learning in Financial Services*' provides a comprehensive overview of current applications of machine learning to market risk management and is recommended reading.

Trading in financial markets inherently involves the risk that the model being used for trading is false, incomplete, or is no longer valid. This area is generally known as model risk management. Machine learning is particularly suited to stress testing market models to determine inadvertent or emerging risk in trading behaviour. Woodall (2017) describes a variety of current use cases of machine learning for model validation, including the French investment firm Nataxis which at the time of writing was running over 3 million simulations a night using unsupervised learning to establish new patterns of connection between assets and investigating further any simulations that emerged from the testing that showed 'errant' patterns compared to average estimates. Woodall also notes how Nomura use machine learning to monitor trading within the firm to verify that unsuitable assets are not being used in trading models. An interesting current application of model risk management is the firm *yields.io* which provides real time model monitoring, model testing for deviations, and model validation, all driven by AI and machine learning techniques.

Another area of focus within the category of market risk of importance to large trading firms is understanding the impact of their trading on market pricing. Day (2017) explores how large trading firms are using AI, and particularly clustering techniques, to avoid the costs of trying to trade into and out of large positions in illiquid markets. He provides a quote from Capital Fund Management, one of the largest hedge funds in France with \$11 billion under management, claiming that up to two thirds of their profit from trades can be lost due to market impact costs. Machine learning techniques significantly help with this issue by identifying connections between assets that are not easily observable and thus allow entering desired positions through a series of related assets rather than taking a large position in a single asset. Cluster analysis particularly helps in this regard (Calvalcante et al, 2016), as can deep learning models (Heaton et al, 2017).

A direction towards pure AI involves reinforcement learning, where market trading algorithms are embedded with an ability to learn from market reactions to trades and thus adapt future trading to take account of how their trading will impact on market prices (Hendricks and Wilcox, 2014). Another interesting direction is proposed in Chandrinou et al (2018), based on tests using foreign exchange market trading data, where a combination of neural networks and decision tree techniques are used to provide real-time warnings to traders of changes in underlying trading patterns while trading. Wu and Olsen (2017; Chapter 5) also examine the use of machine learning to provide warning signals to traders and demonstrate the benefits of SVM as a suitable technique.

### **3.3 Application to operational risk**

Operational risk management entails the firm seeking to identify the risk of direct or indirect financial loss emanating from a host of potential operational breakdowns (Moosa, 2007). These risks can be internal to the institutions (e.g., inadequate or failed internal processes, people, and systems), or from external events (e.g., frauds, vulnerable computer systems, a failure in controls, operational error, a procedure that has been neglected, or a natural disaster). Given the increase in both quantity, variety, and complexity, of operational risk exposures, especially for financial firms, this has introduced a path towards artificial intelligence and machine learning based solutions (Choi et al, 2017).

From the lens of a risk management process, artificial intelligence can assist institutions at various stages ranging from identifying risk exposure, measuring, estimating, and assessing its effects (Sanford and Moosa, 2015). It can also help in opting for an appropriate risk mitigation strategy and finding instruments that can facilitate shifting or trading risk. Thus, use of AI techniques for operational risk management, which started with trying to prevent external losses such as credit card frauds, has now expanded to new areas involving analyzing extensive documentations and performing repetitive processes, as well as detection of money laundering that requires analysis of large datasets.

Taking the example of financial fraud: banks attempt to control this through evaluating the best ways to protect their systems, their data, and ultimately their clients. AI's ability to introduce better process automation can accelerate the pace of routine tasks, minimize human error, process

unstructured data to screen out relevant content or negative news, and determine individuals' connectedness to evaluate risky clients and networks. This same network analysis can also be used to monitor employees and traders. Clustering and classification techniques can be used to establish behavior-based trader profiles, where combinations of trade data, electronic and voice communications records enable banks to observe emerging patterns of behavior to predict latent risks and detect links between employees. It also enables banks to generate and prioritize alerts based on types of suspicious activity and the level of risks involved. Ngai et al (2011) provide an excellent overview of the core AI techniques used in the financial fraud detection, and note the main techniques applied as being decision trees and neural networks.

As a practical application, five of the biggest banks in the Nordics have recently joined together to establish a joint anti-money laundering infrastructure —known as Nordic KYC Utility. An AI based infrastructure will help comply with regulations and requirements related to KYC (Know Your Customer) regulations and avoid the imposition of fines by regulators. Similarly, HSBC is introducing AI technology developed by data analytics firm Quantexa to monitor their anti-money laundering processes. There are practical efforts in fraud prevention too: a joint venture of Royal Bank of Scotland and Vocalink in the UK is creating a machine learning system to scan transactions by small and large business customers to identify and circumvent false invoices and potential instances of fraud. A study by Colladon and Remondi (2017) using real data from 33,000 transactions of an Italian factoring firm shows the effectiveness of such analysis in fraud detection (see also Demetis, 2018).

### **3.4 Application to RegTech**

Compliance with risk management regulations is a vital function for financial firms, especially post the financial crisis. While risk management professionals often seek to draw a line between what they do and the often-bureaucratic necessity of regulatory compliance, the two are inextricably linked as they both relate to the overall firm systems for managing risk. To that extent compliance is perhaps best linked to enterprise risk management, although it touches specifically on each of the risk functions of credit, market, and operational risk.

RegTech, an analogous area to FinTech focused on compliance, has thus arisen to assist firms in the increasing demands of meeting compliance. In this area AI and machine learning is starting to play a significant role driven by the sheer volume of data that needs to be assessed as well as the non-conventional nature of this data. While much of the role of AI and machine learning in RegTech relates to topics already discussed in prior sections, the key advantage of machine learning in a pure RegTech sense is the ability for continuous monitoring of firm activities. Arner et al (2016) note this ability for real-time insights and therefore avoiding compliance breaches rather than having to deal with the consequences of breaches after they have occurred. Other advantages noted are the ability to free up regulatory capital due to the better monitoring, as well as automation reducing some of the estimated \$70 billion that major financial institutions spend on compliance each year.

A key player in this field is IBM following their acqui-hire purchase of Promontory (a 600-staff RegTech startup), and they now offer a range of AI-driven solutions for reducing RegTech costs, demonstrating the widespread industry interest in the area not just confined to startups. One example of what they do is real-time voice conversation analysis to ensure compliance through a combination of their existing IBM Watson AI expertise and the Promontory domain-specific expertise. This involves translating voice to text and then classifying using natural language processing the conversations into categories that can highlight non-compliance. Other applications of machine learning include the automatic reading and interpretation of the implications of regulatory documentation using natural language processing, as a London-based startup Waymark currently offers to financial institutions.

#### **4 The challenges and future of AI and machine learning for risk management**

There are some significant practical issues that need to be addressed before AI and machine learning techniques for risk management can claim its full potential. The most important of these is the availability of suitable data. Although machine learning packages for Python and R can easily read all types of data from Excel to SQL and can perform natural language processing and process images, the speed with which machine learning solutions have been proposed has not kept pace with firms abilities to suitably organize the internal data they have access to. Data is often held in separate silos across departments, perhaps on different systems, and perhaps with internal political

and regulatory issues restricting sharing of data. Important data might not even be recorded as data but rather kept as informal knowledge of the firm.

Another issue is the availability of skilled staff to implement these new techniques. A survey of the top 1,000 firms in the US on AI implementation in their firms found that their biggest concern in the implementation of AI was the readiness and ability of staff to understand and work with these new solutions (Wilson et al, 2017). Training a skilled cohort of staff is something that will take time, although Goldman Sachs, among other firms, has recently attempted to bypass this problem by developing a campus with space for 7,000 workers in India where prevalence of these skills is more present.

There are also practical issues over how accurate machine learning solutions actually are. The range of testing approaches available within machine learning is growing rapidly, and that is a good thing, but it is also driven by the evident limitations of the previous methods and the need to overcome those limitations. This suggests that firms can not simply ‘apply’ a machine learning risk management solution, but rather that it is a continuous process requiring constant evaluation of whether their particular machine learning solution is currently considered best practice. When it comes to AI, where there is some or full automation of process from data gathering to decision making, the need for human oversight will become even more pressing. The case of Knight Capital from back in 2012 serves to illustrate this risk, with their stock trading automation through algorithms resulting in a loss of \$440 million in the space of just 45 minutes<sup>5</sup>. As AI starts to automate lending and credit risk decisions it will be imperative to ensure that such risks can be controlled before handing over control.

This last point serves to introduce the final major issue around transparency and ethics which AI-driven solutions need to further address. Transparency is especially an issue for the increasingly popular deep learning method of machine learning as the models work on hidden layers between the input data and the output decision. A black box system of this type of not conducive to effective risk oversight and can cause regulatory compliance issues especially around demonstrating model validity. A more hypothetical issue related to this is that models used by different firms might converge on similar optimums for trading causing systematic risk as well as loss to the firm. There are also broader ethical issues. For example, discrimination against race in lending decisions is

---

<sup>5</sup> <http://fortune.com/2012/08/02/why-knight-lost-440-million-in-45-minutes/>

widely legislated against, as are increasingly discrimination based on gender and sexuality. Normally these restrictions are incorporated as hard-coded rules in AI and machine learning techniques concerning credit risk and lending decisions. However as deep learning becomes more popular it becomes much harder to police that the model is not inadvertently making decisions based on indirect proxies for these red lines. This is especially the case as more and more atypical data is incorporated in risk management. These ethical aspects are only likely to increase in significance over time.

Leaving these important issues aside for now, it is worth considering the future role for AI and machine learning in risk management from a more positive angle. One obvious conclusion is that the time-consuming and costly nature of risk management will diminish significantly. For example, BBVA, the second largest bank in Spain, has 8,000 of their 137,000 staff working in compliance. They are investing heavily in AI solutions to reduce this compliance cost base. The ability of AI and machine learning to automate repetitive tasks and to organise, retrieve, and cluster non-conventional data such as documents is naturally going to confer cost benefits on firms that move more into this area.

AI will also increasingly deliver accurate real-time information on all types of risks being taken by the firm. As data organisation becomes more orientated towards use by AI, real-time advice will become a pervasive presence. The following step from real-time knowledge of risks being taken is preemptive notice of risks. To some extent this is the holy grail of an AI-driven risk management system – to be able to accurately know in advance firm risks, be they market, operational, or credit risk. The techniques of machine learning offer this ability in the way that traditional statistical techniques could never hope to. Thinking even further ahead there is no technological impediment to a truly AI risk management system that will automatically intervene to prevent unwarranted risks, to immediately unwind dangerous exposures, and to dynamically adjust the risk appetite of the firm based on the system's estimate of the broader risk environment. Although that then will present its own risks which will in turn have to be managed, keeping risk management professionals happily employed (albeit in a fast-changing environment) for the foreseeable future.

## References

- Altman, Edward I. "Financial ratios, discriminant analysis and the prediction of corporate bankruptcy." *The Journal of Finance* 23:4 (1968): 589-609.
- Altman, Edward I., Giancarlo Marco, and Franco Varetto. "Corporate distress diagnosis: Comparisons using linear discriminant analysis and neural networks (the Italian experience)." *Journal of Banking & Finance* 18:3 (1994): 505-529.
- Arner, Douglas W., Janos Barberis, and Ross P. Buckley. "The Emergence of RegTech 2.0: From Know Your Customer to Know Your Data." *Journal of Financial Transformation* 44 (2016): 79-86.
- Cavalcante, Rodolfo C., Rodrigo C. Brasileiro, Victor LF Souza, Jarley P. Nobrega, and Adriano Oliveira. "Computational intelligence and financial markets: A survey and future directions." *Expert Systems with Applications* 55 (2016): 194-211.
- Chandrinou, Spyros K., Georgios Sakkas, and Nikos D. Lagaros. "AIRMS: A risk management tool using machine learning." *Expert Systems with Applications* 105 (2018): 34-48.
- Choi, Tsan-Ming, Hing Kai Chan, and Xiaohang Yue. "Recent development in big data analytics for business operations and risk management." *IEEE transactions on cybernetics* 47:1 (2017): 81-92.
- Colladon, Andrea Fronzetti, and Elisa Remondi. "Using social network analysis to prevent money laundering." *Expert Systems with Applications* 67 (2017): 49-58.
- Day, Sebastian. "Quants turn to machine learning to model market impact", *Risk* 2017 (April)
- Demetis, Dionysios S. "Fighting money laundering with technology: A case study of Bank X in the UK." *Decision Support Systems* 105 (2018): 96-107.
- Domingos, Pedro. *The Master Algorithm*. New York: Basic Books, 2015.
- Figini, Silvia, Federico Bonelli, and Emanuele Giovannini. "Solvency prediction for small and medium enterprises in banking." *Decision Support Systems* 102 (2017): 91-97.
- Financial Stability Board. *Artificial Intelligence and Machine Learning in Financial Services*. Basel: Financial Stability Board, 2017
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. New York: Springer, 2009.



- Heaton, J. B., N. G. Polson, and Jan Hendrik Witte. "Deep learning for finance: deep portfolios." *Applied Stochastic Models in Business and Industry* 33:1 (2017): 3-12.
- Hendricks, Dieter, and Diane Wilcox. "A reinforcement learning extension to the Almgren-Chriss framework for optimal trade execution." In *Computational Intelligence for Financial Engineering & Economics*, IEEE, 2014: 457-464.
- Hu, Yong, Xiangzhou Zhang, Bin Feng, Kang Xie, and Mei Liu. "itrade: A mobile data-driven stock trading system with concept drift adaptation." *International Journal of Data Warehousing and Mining* 11:1 (2015): 66-83.
- Khandani, Amir E., Adlar J. Kim, and Andrew W. Lo. "Consumer credit-risk models via machine-learning algorithms." *Journal of Banking & Finance* 34:11 (2010): 2767-2787.
- Kumar, P Phani. *Machine Learning for Model Development in Market Risk*. New York: GARP Institute, 2018
- Moosa, Imad A. *Operational risk management*. New York: Palgrave Macmillan, 2007.
- Nazemi, Abdolreza, Konstantin Heidenreich, and Frank J. Fabozzi. "Improving Corporate Bond Recovery Rate Prediction Using Multi-Factor Support Vector Regressions." *European Journal of Operational Research* 271:2 (2018): 664-675.
- Ngai, E. W. T., Yong Hu, Y. H. Wong, Yijun Chen, and Xin Sun. "The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature." *Decision Support Systems* 50:3 (2011): 559-569.
- Olson, David L., and Desheng Dash Wu. *Enterprise Risk Management in Finance*. New York: Springer, 2015.
- Sanford, Andrew, and Imad Moosa. "Operational risk modelling and organizational learning in structured finance operations: a Bayesian network approach." *Journal of the Operational Research Society* 66:1 (2015): 86-115.
- Son, Youngdoo, Hyeongmin Byun, and Jaewook Lee. "Nonparametric machine learning models for predicting the credit default swaps: An empirical study." *Expert Systems with Applications* 58 (2016): 210-220.
- van Liebergen, Bart. "Machine learning: A revolution in risk management and compliance?" *Journal of Financial Transformation* 45 (2017): 60-67.



Wilson, H. James, Paul Daugherty, and Nicola Bianzino. "The jobs that artificial intelligence will create." *MIT Sloan Management Review* 58:4 (2017): 14-16.

Woodall, Louie. "Model risk managers eye benefits of machine learning." *Risk* 2017 (April).

Zhong, Xiao, and David Enke. "Forecasting daily stock market return using dimensionality reduction." *Expert Systems with Applications* 67 (2017): 126-139.

.

Figure 1: Categories of machine learning technique

		Linear methods	Non-linear methods
<b>Problem type</b>	<b>Supervised</b>		
	Regression	<ul style="list-style-type: none"> <li>• Principal components</li> <li>• Ridge</li> <li>• Partial least squares</li> <li>• LASSO</li> </ul>	Penalized regression: <ul style="list-style-type: none"> <li>• LASSO</li> <li>• LARS</li> <li>• elastic nets</li> </ul> Neural networks and deep learning
	Classification	Support vector machines	Decision trees: <ul style="list-style-type: none"> <li>• classification trees</li> <li>• regression trees</li> <li>• random forests</li> </ul> Support vector machines Deep learning
	<b>Unsupervised</b>		
	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.	

Source: Table 1 from van Liebergen (2017)