Using Machine Learning in Electronic Trading

Final Report

Professors Fethi Rabhi and Nikolay Mehandjiev[1]

1 Introduction

The area of electronic trading has seen dramatic changes in recent years. In particular, trading in financial markets has become a global activity because of the recent technological developments that have facilitated the instantaneous exchange of information, securities and funds worldwide. The economic factors behind this push are transparency, cost, risk management and the potential for anonymity. Electronic trading not only improves transparency of prevailing prices in a market but also provides more information such as the depth of a market which indicates the potential supply and demand away from the current market price. At the same time, the exponential increase in computing power and data storage during the last decade has resulted in the rapid development of machine learning and data mining with diverse applications in economics, finance, science, engineering, and technology. In the finance area, machine learning models have elicited considerable attention from many researchers because of their predictive power.

The motivation of this study is to understand how machine leaning techniques are being applied in the area of electronic trading from a broad angle encompassing both academic and industry perspectives. Although the interest of the academic community for machine learning techniques is growing, there is still a large gap with professional practice [7]. One of the reasons is that different research communities understand machine learning in different ways. To make matters worse, this concept is also used interchangeably with that of Artificial Intelligence or AI. This has necessitated a study to analyse the state of adoption of machine learning in the electronic trading community.

In the rest of this report, Section 2 gives some background and existing literature, also explaining the research methodology used in our study. Sections 3-5 discuss the results of conducting semi-structured interviews with industry experts. Conclusions are presented in Section 6.

2 Background and Existing Literature

2.1 What are machine learning and electronic trading?

A recent survey conducted by the Bank of England and the Financial Conduct Authority in October 2019 [2] defines machine learning as "a methodology in which programs fit a model or recognize

---

[1] Disclosures: This paper has been sponsored by the Plato Partnership. However, the views expressed in the paper are the views of the authors and do not necessarily reflect the views of the Plato Partnership. Fethi Rabhi is a Professor in Software Engineering at the University of New South Wales and Nikolay Mehandjiev is a Professor in Information Systems at the University of Manchester.

patterns from data, without explicitly programmed and with limited or no human intervention". It considers that machine learning constitutes an improvement rather than a fundamental change from statistical methods. Machine learning can also work hand in hand with any type of automated optimization techniques that allow a computer to "twiddle the knobs", i.e. to search the space of possible combinations of parameter values, to find good settings [3].

Electronic trading is associated with the area of financial markets so it is mainly concerned with the automation of traditional trading activities in the finance domain and their associated processes. Most literature is primarily concerned with order-driven (auction) markets in which trading occurs in different ways such as single-priced auction, continuous rule-based two-sided auction and crossing networks. In this type of markets, trader orders for any particular security are centralized in a Centralised Limit Order Book (CLOB) where buyers are seeking the lowest price and sellers are seeking the highest price. An order-driven market uses order precedence rules to match a buyer to sellers and trade pricing rules to price the resulting trades. As trading has become increasingly fragmented, there are now multiple order books competing against each other, with variation in rules and pricing across the different books. In these cases, integration is achieved through smart order routers.

In such markets, trading decisions have also become more complex as they involve the cooperation of different participants and their systems interacting across wide geographical boundaries and different time zones. For these reasons, machine learning techniques can be applied in different parts of the electronic trading lifecycle to increase efficiency and accuracy of such decisions. Many researchers also believe that machine learning has the potential to replace theoretically derived models with data driven models [4].

2.2 How does academia envision the use of machine learning in electronic trading?

Analysing data patterns in financial markets represents the largest body of research that was uncovered during the literature study. Indeed identifying patterns such as stock price movements has been a major challenge for both academics and practitioners for decades. To address the limitations of traditional econometric models, machine learning models have become popular for their ability to process complex, imprecise, and voluminous data [10]. In addition, these methods enable the use of different types of data (qualitative and quantitative) and are not subject to rigid assumptions such as those imposed on econometric models. They have the freedom to incorporate fundamental and technical analysis into a forecasting model and can be adapted to different market conditions.

The rise in the use of machine learning techniques has also triggered a lot of interest in developing innovative ways of gathering and processing data from multiple sources in order to produce good quality datasets. The motivation behind these efforts is that despite the availability of numerous datasets, they are not quite ready to use, needing further processing to deal with their massive size and high dimensionality. Related work differentiates between the analysis of structured and unstructured datasets as they present different characteristics. It is recognized that the direct use of low-level market data in any type of model is not recommended because of heterogeneity, noise accumulation, spurious correlations, and incidental endogeneity [7]. Despite the fact that machine learning techniques can deal with more features than traditional econometrics models, adding more predictors is not a guarantee for increasing the performance of the analysis [7].

In addition, there has been extensive research work done in textual information representations that enable the exploration of relationships between news content and financial markets. More recent work in textual analysis has focused on social media data. Any research work in the area of sentiment analysis is also relevant to the area of financial trading as sentiment affects markets. The use of social media to implement collaborative trading strategies which have the potential to skew markets in the short term has been observed within the recent case of Reddit[2] users (mostly retail investors) conspiring to manipulate the price of a small number of companies including GameStop.

Finally, as decision-making has become very complex at trading desks, the increasing use of algorithms is being seen as essential to determine impact on market liquidity and volatility [6]. There are many examples of using machine learning in areas such as solving complex optimization problems, order placement, intraday automated trading and dealing with operational efficiency problems during high frequency trading. Despite these examples, academic research work is often lagging behind industrial innovations in this area, as evidenced by the proliferation of information about new approaches and applications coming from industry outlets.

2.3 Practice of machine learning in the wider context of Financial Services

The literature survey has revealed a mismatch between the focal areas of academic researchers, which generally reflect classic academic disciplines divisions. For example, most studies in computer science tend to focus on the computerized analysis of information (both structured and unstructured) for the purpose of forecasting stock prices. There is also a wide gap between academic finance and "professional" finance when it comes to analysing big datasets with machine learning methods and putting these models into operation within a realistic trading environment. The literature survey shows that academic literature is not sufficiently addressing some of the major trends that have been shaking this sector in recent years:

1. **De-verticalization and move towards a service economy**: important changes in the regulatory frameworks (particularly the MiFID regulations in Europe) have disrupted the vertical separation between "sell side" (e.g. investment banking that have been enjoying privileged access to market liquidity on centralized exchanges) and "buy-side" fund managers [3]. More stringent compliance requirements have changed the relationships between equity trading businesses and their brokers. Buy-side companies can now pick-and choose their technology components as well as their trading venues and therefore avoid paying for the functionality previously provided by investment banks [13]. Consequently, the Fintech sector has witnessed a rise in Small and Medium-sized Enterprises (SMEs) offering software services and component technologies that can perform various functions in the trading cycle.
2. **New alternative trading venues:** there has been significant growth in the provision of new alternatives to existing trading venues. This includes "dark pools" provided by off-exchange trading venues, called Alternative Trading Systems (ATSs) in the US and Multilateral Trading Facilities (MTFs) in Europe. They allow large blocks of shares to be traded with a higher degree of anonymity. Despite of the global increase in trading venues, 33% of traders in a recent survey still mention liquidity availability as an important issue in 2020 [9].
3. **Extending information reach:** there has been a huge rise in the availability, variety and volume of data associated with the different models used during the trading cycle, as well as a diversification in the types of instruments that can be traded. In particular, complex derivatives and structured instruments such as exchange traded funds (ETFs) require additional information to be processed to enable complex predictions to be made. At the same time, the increasing use of sophisticated machine learning techniques (particularly deep learning) in pre-trade and post-trade analytics require more information to be processed to enable faster and more accurate decision-making. A recent survey is listing "analysis of previously inaccessible data" which is "reshaping trading strategies" as a key driver in this area [9].
4. **Faster speed and more automation:** the trend in the use of high-speed, high bandwidth, adaptive technologies in automated systems interacting with a wide range of trading venues is set to

[2] Kiran Stacey. GameStop mania: why Reddit traders are unlikely to face prosecution. Financial Times, last retrieved 31st Jan 2021. Available from https://on.ft.com/3tejsMc

continue over the next decade. Although high-frequency traders only represent a small proportion of trading parties, they are responsible for a majority of the equity trading volume. As most existing models and systems have been designed to operate with static data, there is now a need for new cost-effective solutions that are able to acquire, process and analyse real-time data. In addition, the need to integrate several systems (possibly from different vendors) has revealed many issues related to workflow design and efficiency. Finally, the tendency to use machine learning techniques to assist in selecting the most appropriate trading strategy based on a user's parameters is set to continue as only 27% of firms to date have tried "AI-powered algos" according to a recent survey [9].

5. **Need for better risk management:** all of the above trends and in particular the need to integrate multiple technologies and models exacerbate the potential for massive risk throughout the trading cycle. In addition, the risks associated with automated trading are poorly understood especially the quantification of risk which requires a new risk management culture to be established in financial institutions and regulators [3].

2.4 Research methodology

Further investigation was needed to explore the following themes from a practitioner's perspective:

(1) What is machine learning and how is it being applied in electronic trading?
(2) What are the key management and technical aspects of machine learning in electronic trading?
(3) What are the gaps and areas of future research and actions?

During 2020, 15 exploratory semi-structured interviews were conducted with industry experts from UK, Australia, USA and Germany. Data from the interviews was qualitatively analysed and the results are presented in this report as follows:

- Section 3: Machine Learning Adoption in Electronic Trading
- Section 4: Machine Learning's Place in Corporate Strategy and Operations
- Section 5: "Under the hood" of Machine Learning

3 Machine Learning Adoption in Electronic Trading

3.1 Diversity of understanding of Machine Learning and Electronic Trading

When asked about their understanding of machine learning, a large number of interviewees used informal definitions like "when there is no explicit algorithm", the use of big data to make decisions, extracting knowledge or being able to explain complex problems to a non-technical audience. Interestingly, some defined machine learning by contrasting it with AI and stressing that the two are different. Others defined machine learning according to one specific technique such as neural networks or text analysis. The majority presented machine learning as an extension of statistical analysis of data, just another way of addressing the limitations of existing techniques for extracting value from data. Unlike classical statistics, it can deal with non-parametric models, discover nonlinear relationships and handle a high number of parameters. Some interviewees also mentioned that some advanced machine learning techniques are not static i.e. they are able to change and adapt over time.

Similarly, electronic trading is understood in very general terms by some of the interviewees like using an algorithm or replacing human intervention by an algorithm during trading. The majority do focus on the use of computers and diverse ways for automating the process of trading. Many have made references to various places in which such automation occurs from determining order timing and size decisions to order execution. Some have pointed out particular areas in the trading cycle that are challenging for automation such as trade strategy selection.

3.2 Machine learning use cases in electronic trading

There were many opinions related to how Machine Learning Techniques can be used in the area of electronic trading, summarised in the table below.

Area	Sub-area	Selected quotes
In general	Anything requiring speed & precision	“doing this at precision and speed humans cannot do (we cannot press the button fast enough) - trading at shorter time frames”
	Can be used in many places	“ML used for small parts, not for overall decision-making” “No need to pinpoint problem” “Depends on labelling you can do”
	Where transparency is needed	“think through everything that is happening in the model, do not use blindly”
Supporting specific areas of trading process	Order generation & submission	“in order placement and order volume prediction (eg placing orders depending how volume of the day changes).” “order timing, size and routing decisions use ML”
	Trading signal generation	“at high frequency or structured arbitrage trading, etc -ML to X ray markets by pinging single shares”
	Trading strategy selection	“supports a buy-side decision maker (trader) to pick a specific broker algorithm for one specific stock or trade that has come in from a Portfolio Manager”
	3.2.1 Client Flow Management	“managing client flows”
	Venue Selection	“Focus of ML applications would be, for example, venue selection and main approach would be semi-automated”
	Order execution	“order routing and execution - where all the hard work is done” “Execution using machine learning: a part of "order working", making a decision about "crossing the spread or not"”
Supporting analytics	Extracting patterns in data	“pattern detection (structured and unstructured data (company reports, news)” “pre-trade analytics: pattern recognition and predictive techniques”
	Modelling complex systems	“The problem is conceptualised by studies in complex systems”
	Asset-specific	“The purposes are asset-specific”
	Evaluation	“You can measure performance, I guess, at all these different levels, too. And so, trying to optimize the performance is just an optimization problem and you can solve that with ML and a lot of cases”

As can be seen from the table, the use of machine learning varies enormously in the trading cycle. It can be directly applied to improve automated decision-making and trading activities or to support various types of analytics for any type of trading (automated or human-based).

3.3 Perceived benefits

Economics teaches us that automation can improve productivity by increasing the *scale* of activity, the *scope* of activity and through *learning* used to optimise operations and decisions. The opinions of experts covered all these three angles in relation to machine learning, illustrating the fact that machine learning is perfectly placed to improve productivity of electronic trading operations.

3.3.1 Machine learning helps companies increase the scale of what they do.

Machine learning was perceived as helping companies to increase the scale of their operations by enabling the design of core trading algorithms in an efficient manner, which can then scale outwards to other financial instruments or markets / exchanges. Furthermore, this can be done with improved accuracy of decisions, demonstrated by randomised testing.

The scale and the variety of data which can be processed is also increased, with machine learning being able to optimise multiple parameters at once on huge data sets, enabling real-time analysis of data generated by contemporary markets.

Machine Learning can also identify patterns not only more accurately but also quicker compared to using conventional observations, which can bring competitive advantage. The costs are also reduced, allowing decisions to be based on larger sets of finer-grained observations.

Overall, machine learning helps increase the scale of operations by speeding up decision-making processes in a repeatable and scalable manner.

3.3.2 Machine learning helps companies expand the scope of their electronic trading.

Machine learning was reported to help address problems which are difficult to formalise and hence to automate, thus extending the scope of what an organisation can try. This helps improve the breadth in the diversity of electronic trading instruments, for example application of decision-making algorithms at different levels of granularity that were previously impossible.

Another area which is impossible for human experts to tackle without the assistance provided by machine learning is extracting complex patterns from huge volumes of data. It can help to discover unexpected interdependencies between variables which will not be tested by human analysts.

3.3.3 Machine learning helps companies to improve by learning from past experiences.

A strong aspect of learning feedback was evident in the majority of responses. For example, one of the best features enabled by machine learning was the ability of an algo-wheel to gather real-time feedback on its performance, and the data thus gathered will be analysed to decide on the extent to which the new learning will be incorporated in core algorithmic code. This mechanism recognised the role of human experts in committing the learning to the core operating code, yet the feedback's main enabler is using machine learning itself.

Machine learning can also help avoid some cognitive biases of human learning, where traders remember successes better than failures, leading to them always using the same trading algorithms in a kind of "muscle memory". Objectively evaluating the performance of algorithms in real time can lead to adaptive choice of the trading algorithm optimised for the current context in space and time. To quote one of the interviewees:

> *Overall, the machine learning brings about improved trading performance because it engenders higher decision-making accuracy and delivers more sophisticated decision-making mechanisms, able to learn from real-time feedback based on objective performance parameters.*

The participants also noted that the largest impacts of machine learning benefits are seen in upstream decision-making. Going downstream to markets simplifies and speeds up decisions, reducing the decision-making options available and somewhat reducing the impact of machine learning.

3.4 Perceived constraints

The constraints shaping the application of machine learning to electronic trading can be conceptualised using the system diagram below. Machine learning and automated trading (eTrading) work together to deliver good performance (referred to as α) under the control of management and regulator, using data sampled from markets and past orders.

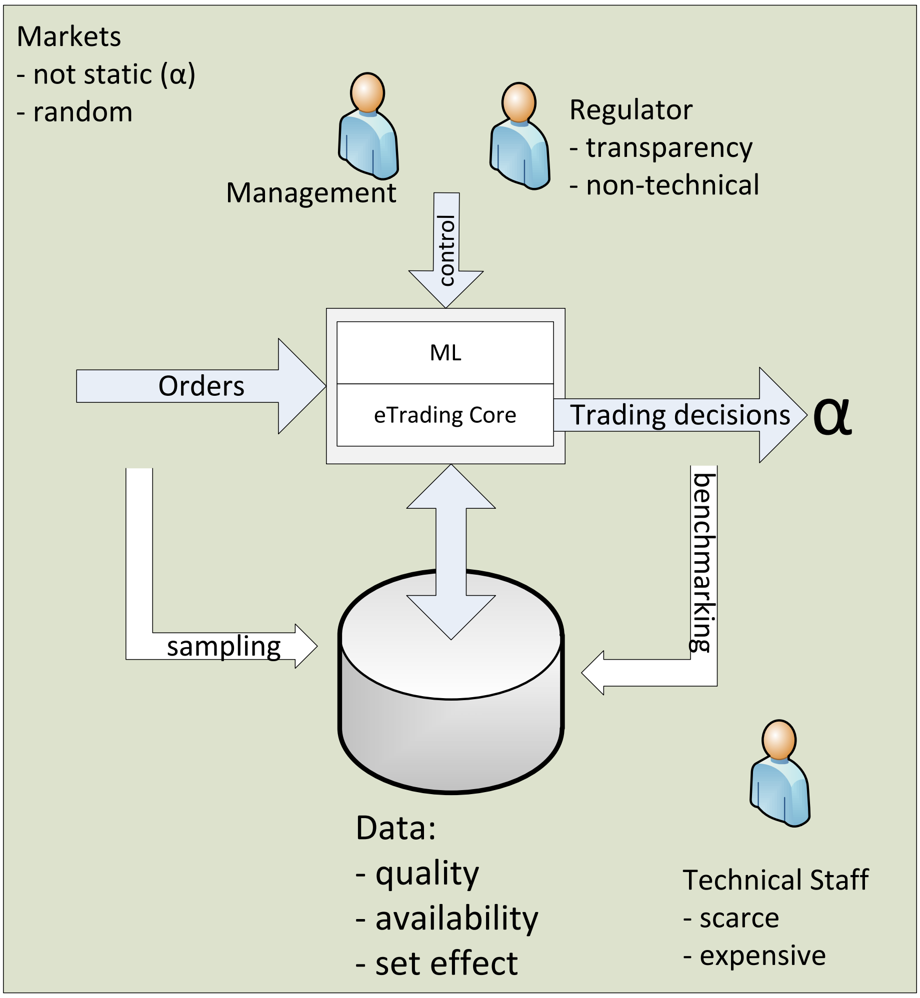

Figure 1. Constraints applicable to an intelligent trading system comprising a Machine Learning component and an eTrading Core component. The system is presented using the IDEF0 conventions with its Inputs (Orders), Outputs (Trading Decisions), Controls (Management and Regulation) and Resources (Data and Technical Staff).

Machine learning is based on two fundamental assumptions – that the observed data is not random, which contradicts the random-walk presumption about markets; and that the observed data will remain static for a sufficiently long period for machine learning to work, after training on sampled data has been completed. The problem with the second assumption is that markets change in both dynamics and nature, as other traders also discover ways to deliver excess returns.

In addition, the available data often lacks the quality necessary to make machine learning effective. Training a machine learning model on a smaller than necessary subset of data brings set effect error and results in over-fitting, where predictions are brittle when applied to out-of-sample data.

Machine learning technology has many inherent characteristics which constrain its wide application to electronic trading: it is perceived too slow to bring value downstream where speed of execution requirements are much tighter, indeed the majority of machine learning applications are upstream in the analysis part. Artificial neural networks (ANNs) are especially slow to train and operate, so our respondents did not report any notable uses for them. Unsupervised learning algorithms were perceived as not very applicable to the problems addressed by electronic trading, in particular clustering was criticised for difficulties in balancing between cluster populations.

Even more critical was the perception of some machine learning techniques as focusing on optimising a few target parameters in an isolated manner, leading to the Goodhart's Law. There machine learning was delivering optimal results on these isolated parameters but failing to achieve the overall aims of gaining the best alpha for the overall trade. Often machine learning models would fail to trade parts of the order which do not bring sufficient gain, only for these to be traded at a loss on the next day.

Another problem was the "black-box" nature of machine learning, which impeded efforts to ensure regulators are convinced that machine learning-driven eTrading complies with regulations. The lack of transparency of machine learning is reported to be confounded by the regulators apparently lacking technical acumen.

Separating machine learning in an "advisory" unit on top of the core eTrading algorithms was seen by some respondents as the answer to this conundrum. This arrangement is seen as ensuring the core eTrading system remains transparent and regulatory compliant, whilst the machine learning advisory unit is only steering the operations of the core eTrading system and cannot through this break any of the regulations abided by the core system. When integrating the machine learning part with the eTrading core, problems were perceived to arise when the trading modalities were split between electronic and human trading.

The final main constraint mentioned was the scarcity and hence of the cost of machine learning experts and also the costs of the infrastructure necessary to run machine learning with the necessary performance.

4 Machine Learning's Place in Corporate Strategy and Operations

4.1 Introduction

The use of algorithms in trading in Europe is regulated by MIFID II, which stipulates that algorithms used in trading should be registered, tested and should include documented circuit breakers. These provisions are expanded for UK-based trading by the Bank of England PRA's Supervisory Statement SS5/18. This statement describes PRA's approach to risk controls, governance and testing of algorithmic trading, covering the following five main areas:

1. Governance (roles and policy)
2. Algorithm approval process
3. Testing and deployment
4. Inventories and documentation
5. Risk management and other functions

In the sections below we structure these areas in three themes emerging from the diversity of responses to our questions regarding governance and validation. Our respondents exhibited different degrees of awareness of Governance and Risk practices in their organisations, yet the need for transparency regarding the complex machine learning algorithms and having to make them understood by non-technical regulatory bodies was universally mentioned in most responses.

4.2 Governance and Strategic Alignment

Our respondents tended to agree with the generally accepted notion that to gain the full benefits of advanced machine learning applications, their introduction and control should be aligned with the business strategy, and a clear business goal should be communicated to the machine learning activity. The need to keep the organisation aligned to the industry standards and regulations was critical as one of the respondents stated that his organisation needed to deal with "regulators in 52 markets and 104 different exchanges with constant changes". The responsibilities are understood to be of the competence of the CEO role and several dedicated committees as described in the section below.

4.2.1 Roles

Our interviewees referred to the following roles tasked with governing electronic trading and hence the application of Machine Learning in electronic trading:

- The CEO of a company should be responsible for the overall alignment of business objectives, machine learning/eTrading activities and regulatory documents. This view is aligned with a presumed top-down change approach driven by a key person.
- The "regulatory" team which covers all the regulatory matters within the company is referred to as the "compliance and legal" team. It was stressed that if this team is to effectively deal with regulating machine learning activities, it needs to be multi-disciplinary including technical experts.
- The "model validation" committee was also mentioned as the one responsible for the formal Algorithms Approval Process covered below.
- The "execution review" committee which would ensure the compliance with the "Best execution" requirements was also mentioned.

Some interviewees also saw the role of a separate specialist team ("ML Lab"), focused on the machine learning strategy, fully independent from the team that built core trading algorithms. This distinction was at the heart of an approach which saw Machine Learning as fast-changing exploratory add-on to a stable and fully validated core eTrading system, which encodes regulatory compliance rules.

4.2.2 Dedicated policy regarding Algorithmic Trading and Machine Learning aligned with Company Strategy

Our respondents tended to perceive algorithmic trading as sufficiently different from machine learning activities to warrant different regulatory policies. Machine learning was seen as a "tool suitable for the job", and so different company goals would pose different requirements to the machine learning operations. Many companies are client-driven, so that the scope of the ML activity is focused on "executing clients' orders at minimal cost."

For the trading-focused companies and companies focused on pricing of instruments (e.g. derivatives market-making firms), machine learning was seen as a supporting capability improving their core activities. It was then considered both small in scale and requiring rapid configuration changes, rendering it unsuitable for the usual regulatory processes. It is also "new" so not yet in scope of company regulatory activities. The proposed configuration was for the machine learning sub-system to steer the core algorithmic trading part of the system in ways which do not break regulatory compliance. Then it would only be necessary to have dedicated regulatory policy regarding the algorithmic trading part of the system.

In contrast, for companies specialising in adding value by developing innovative trading algorithms and providing pre-trading analytics services, machine learning is perceived as more central and so subject to full regulatory compliance policies.

One issue which united different companies according to their representatives was the challenge to explain machine learning models to non-technical regulatory agents, so they can understand the techniques and reasons behind different decisions. This is explored in detail below.

4.3 "X-raying" the black box - Provenance, Explainability and Transparency

The issues of transparency and understanding of machine learning models stem from the black-box nature of such systems. Our interviewees were fully aware of the challenges this poses to the uptake and regulation of machine learning as applied to the algorithmic trading part of electronic trading. These problems were perceived to affect three of the areas covered in the BoE PRA's approach: Algorithms approval processes; risk management, and inventory/documentation issues.

4.3.1 Algorithms Approval Process

Some interviewees reported that their process for approving any algorithm was "heavily governed and formalised", and based on testing the performance of the algorithm. Certain aspects of these processes were perceived as slow to "catch up" with the specific aspects of machine learning models:

(a) Machine learning is based on data, indeed its operation and results are based on the data used for training and then testing. The size and scope of the data sample used (e.g. one day of trading) is perceived as being sufficient for conventional algorithms yet too small for machine learning approaches. This small sample brings significant distortions of the results reflecting the specific characteristics of that particular day of training.

(b) The statistical complexity of machine learning instruments was highlighted as one of the main barriers to the operation of regulatory approval processes. Indeed, the machine learning models are reportedly perceived as black-boxes not only by risk managers from non-technical background but even by the computer-science specialists responsible for their implementation.

4.3.2 Inventory and documentation

Recording the features and documenting the testing conditions for the machine learning algorithms and models were deemed important for the success of machine learning's application to electronic trading. For example, the difficulties regulators faced in understanding the limitations of testing data-intensive models with restricted data sets were permeated by failure to document these limitations of the approval process. Respondents highlighted that "sell-side lacks incentives for sufficient validation", with no standards in place which would force them to include "out-of-sample performance" in the documentation made available to buy-side users.

This lack of performance and scope-of-testing data creates significant over-fitting risks, and difficulties in comparing the performance of two alternative implementations of a machine learning model, or of two different machine learning models when selecting the best tool for the job at hand.

4.3.3 Risk managers and others should understand trading algorithms and risks

The problems with trying to manage the risk and ensure regulatory compliance of trading systems based on machine learning stem not only from their data-driven nature and technical complexity, but also from their black-box nature, where the "learning" creates non-transparent dependencies between inputs and outputs. Even if the outputs are deemed correct on observation, this lack of transparency makes risk managers unsure that the tool will continue to produce correct outputs for different combination of input parameters.

The problem of trying to explain black-box operations is endemic as artificial intelligence is adopted widely throughout our society. It is, however, particularly pertinent to managing the risk of machine learning in electronic trading because of the speed and monetary impact of decisions driven by this technology in trading.

Companies attempt to address this problem by investing in explainable solutions, and by ensuring the risk management boards have inter-professional membership with stakeholders from the whole organisation including IT support.

4.4 Validation and safeguarding the operations

Our respondents confirmed that thorough generic validation rules are applied across all algorithms and machine learning applications used for electronic trading in their organisations. Testing was seen as taking place at four levels:

1. Is the machine learning model using the right features (especially when these are selected by the model itself)?
2. Is the model correct?
3. Does the model work in determining the correct trading strategy when back-tested?
4. Is the model working in practice on new data?

Some of these are done before the approval of the algorithms, others take part during the operation of the models, and are complemented by safeguards in case the models "misbehave".

4.4.1 Testing and Deployment of Algorithms

Validating machine learning models and electronic testing algorithms before deployment includes:

- Demonstrating the consistent operation of the algorithm across a number of data samples and entire data sets; using a classic experimental design by instrument, by sector and by market.
- Ensuring the right data granularity is used, and that the data is with sufficient volume for statistical validity of results. This includes removing outliers which can distort results, and exchange-specific "noise" data, for example test data or "warm-up" data at the start of trading.
- Testing the applicability and accuracy of trading on various venues e.g. comparing random vs routed selection of venues.
- Ensuring regulatory compliance, e.g. best execution, how wide a spread can be etc. requirements.
- Ensuring appropriate fall-back behaviour when exceptions are encountered, to enable "human-in-the-loop" safeguards during operation.

4.4.2 Safeguards and Human in the Loop

Validation does not finish once a tool is put into operation. Indeed, the operation of the algorithm needs to be monitored to ensure it works reliably, and that it reports bad data correctly when this is encountered. Thus, we can ensure that any mismatch in behaviour is not due to data collection quality issues, which have resulted in the model being trained wrongly.

Whilst back-testing is usually sufficient for regulators, our respondents firmly believed in the importance of ongoing monitoring of operations to detect overfitting to restricted data sample during training and testing, changes in market dynamics caused by other traders searching for "alpha" and other related constraints covered in detail in Section 3.4.

Some challenges with monitoring practices include the lack of manual (qualitative) review of results. This may lead to instantiations of Goodhart's law, where the software produces excellent results in terms of optimising one parameter at the expense of the actual trading goals. Appropriate benchmarks should be chosen to measure outcomes against the goals of trading.

The operation within a changing market context motivates the need for human-in-the-loop safeguards, indeed many of our respondents had these implemented as a matter of course, since the outcome of Machine Learning models, even when acted upon by electronic trading, would still end up brought to the attention of human traders on their trading blotters. Human traders monitor for strange patterns (for example all orders are submitted to the same algo). Explict human-in-the-loop monitoring was also deemed necessary, especially when testing new machine learning models. Systems in place often combined:

(a) explicit automatic safeguards, where the system stops trading if unusual conditions or outputs are encountered) to avoid making damaging wrong choices – that was the case during unusual market conditions at the start of the COVID pandemic, or on Brexit day;
(b) human monitoring, where humans monitor a selection of trades depending on the frequency, speed and size of trading;
(c) "dead-man's handle", where the system will realise if a human is not monitoring it and take appropriate action (i.e. alert another human or stop trading).

Machine learning models often attract stricter automatic safeguards than conventional electronic trading algorithms, for example bounding outputs to price predictions to avoid unfortunate combinations of input data resulting in "runaway" outputs, or tighter limits around output actions.

4.4.3 Summary

In summary, machine learning models attract stricter validation procedures and safeguarding measures because of the substantial limitations on the predictability of the outputs produced by machine learning models. These limitations are caused by models' data-driven nature, complex internal mechanisms making their operation black-box, unpredictable and unescapable changes of market dynamics rendering models invalid, and dangers of over-fitting and Goodhard's Law.

5 "Under the hood" of Machine Learning

5.1 Machine learning methods of choice for different use cases

When asked about methods of choice, interviewees gave a range of answers ranging from specific classes of machine learning techniques to techniques that were required for machine learning to work (see table below). The answers illustrate the variety of machine learning methods and the difficulty in dividing them within a clean hierarchical taxonomy. Some classes of machine learning methods that were mentioned contain others (e.g. Neural Networks is part of supervised learning) and some classes may use others (eg. Narural Language Processing use Neural Networks). In many cases, different methods can be used together (ensemble methods). The answers also mention many techniques that are not proper machine learning techniques but are nevertheless used extensively alongside machine learning techniques like pre-processing, feature selection and visualization.

Area	Sub-area	Selected quotes
Class of machine learning Techniques	Supervised Learning: a machine learning technique that has been configured with a training dataset	"supervised learning most relevant for eTrading (regressions for market forecasting, classification for fraud detection, market surveillance)" "Range from linear regression to non-linear models, depending on the time horizon, nature of features being used in the model, degree to which features have been engineered/tailored to the trading problem"
	Neural Networks: a supervised machine learning technique based on layers of artificial	"Neural networks are being used (poor for larger data sets, overfitting and on customization)"

Area	Sub-area	Selected quotes
	neurons or nodes that mimic the way the brain operates	
	Advanced regressions: extending classic regression techniques to deal with more complex problems	"Non-linear regressions (eg on order book dynamics - relationship between the depth of orders in the book and liquidity)." "linear models (linear regression, penalized linear regression, logistic regression) very popular because they're easily interpretable mitigating risks"
	Ensemble methods: running various algorithms and then select the best one based on the accuracy score	"multi-level regressions and tree-based ensemble methods to build a model"
	Tree based methods: the use of a technique to create a number of decision trees describing patterns in the data (for example Random Forest)	"Random Forest (best)" "Tree-based methods (risk of overfitting)" "tree-based methods in factor investing"
	Heuristic models: Mostly algorithmic techniques not underpinned by a formal model.	"heuristics characteristic models may not be ML since algos creator decides what they think is sensible and then they fit some parameters but those parameters are not explicitly optimizing for a goal" "order timing and size: mostly heuristic models, ML used to some extent by some brokers but not as much as pre-trade or post-trade analytics"
	Reinforcement learning: a technique where the system is reinforcing parameters and decision paths which lead to positive outcomes.	"Reinforcement learning for path-dependent models of trading where profitability it the goal to be maximised", also "loop-based optimisations"
	Unsupervised learning: identifying patterns in data without a training set	"Unsupervised learning methods: identifying patterns and structures in exchange data that we may miss, or finding relationships we expect to see but in a more scalable, efficient and accurate manner." "Simple Data Clustering"
	NLP: using natural language processing techniques to analyse and interpret textual information	"NLP: for inputs such as news feeds, sentiment analysis etc"
	Other	"Techniques useful for speech recognition, e.g. Hidden Markov models"

Area	Sub-area	Selected quotes
Data preparation, interpretation, presentation & visualisation	Dimensionality reduction and Feature Selection: using techniques to reduce the number of factors characterising a data point, for example Principle Component Analysis (PCA).	“dimensionality reduction in data cleansing” “in a PCA or similar, apply machine learning to detect market events”
	Pre-processing: transforming data to make it suitable for machine learning	“Need data vetting not only fitting” “Data cleaning or data quality” “Training in real-time” “Labelling is most important” “Random allocation of training and testing data in real time”
	User-friendly interface: presenting data and results in a way that makes sense to humans	“Recommendation to traders” “Visualization” “Explainability: Attempts in front-ends to explain components - mostly outputs not model components. "very much scratching the surface of explainability"”

5.2 Data sources, quality issues and costs

A summary of the data used with machine learning methods is illustrated in Figure 2. The vast majority of interviewees indicated that structured market data is mostly used, particularly fine-grained data like orderbook data which originates from exchanges. They indicated that there are also other types of structured data such as prices and market measures which are often derived from orderbook data. Forecasts for a particular type of data can also be made from not only historical data of the same type but also other sources such as macroeconomic data and company data. Surprisingly although unstructured data was mentioned a lot in its diverse forms, interviewees were quick to point out that their value has been quite disappointing when used within machine learning methods. There was some scepticism pointed out about social network data in particular. One interviewee mentioned that the analysis of unstructured data is more suitable for investing and long-term trading strategies than real-time electronic trading.

5.3 Technical support for testing and safeguarding

In addition to good management and governance practices, interviewees stressed the importance of strong technical support for validation and safeguarding activities particularly when deploying machine learning in an operational environment. In addition to the use of conventional IT security mechanisms (e.g. firewalls, duplication, backup system etc.), there are many precautions that can be taken during model design, at the software implementation stage and when deploying the system.

At the model design level, interviewees mentioned several tools and techniques that automate important tasks such as cross validation (segmenting data into similar buckets when training and testing), back testing (using randomized testing and randomized agents whenever necessary) and other model validation techniques based on real output data (e.g. comparing predicted with actual outcomes).

Figure 2: Data sources used by machine learning methods in electronic trading

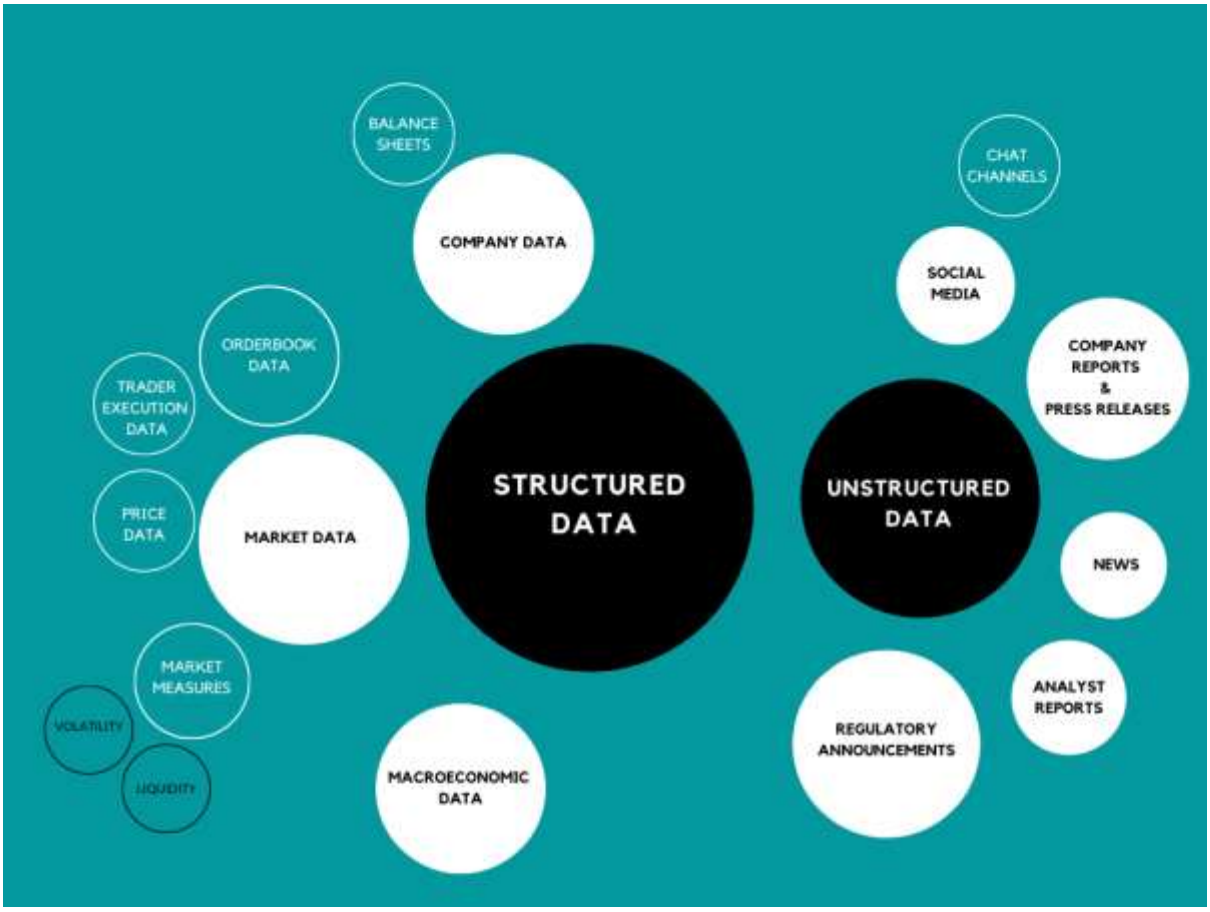

During software development, interviewees stressed the fact that machine learning processes needed to work within the constraints imposed on the fully-compliant and validated core code. One respondent mentioned the need to not only establish safety limits within a machine learning algorithm but also to have independent protection across multiple layers to shield against any individual layer not functioning correctly". Voting or averaging is also mentioned alongside duplication of code to avoid spurious decisions made by a particular software bug.

At deployment time, the software must be gradually deployed on a partial use basis. During operations, various safeguards around the system need to be implemented. Examples include guardrails like caps and collars for controlling trading parameters. Examples include apply limit prices not far away from the arrival price to avoid disrupting the market, uncalibrated parameters should be adjusted based on the level of risk and defining price tolerances and dollar value tolerances and percent tolerances and filters.

Automated alert systems are considered critical. Performance needs to be monitored in real-time and kill switches must be activated to prevent disasters. According to one respondent, there is a need for "somebody to pull the plug" or otherwise fall back to a simpler system to avoid crashes.

6 Conclusions

6.1 New insights into electronic trading information flows

The analysis of interview responses, our previous literature survey [12] and further literature analysis [6] have illustrated the fact that electronic trading systems are highly complex and can be viewed differently depending on a participant's particular view (buy-side, sell-side etc.) as well as business model (agency trading, market making, internalisation etc.). For example, the information flow within a sell-side system to manage customer and proprietary order flow is presented in Figure 3. It will comprise the following components:

- **Order generation**: there are many possible business activities that lead to generating an "order" depending on different business models (buy-side investment, proprietary trading etc.). Such an order will still have many unknown parameters. We consider this component to be some form of upstream decision-making which will specify some values and constraints regarding order size, possible venues, timing etc. that will need to be preserved when the order is further refined downstream (the figure refers to the decisions made as "parent-level decisions").
- **Order execution**: consists of the activities typically conducted by trading desks. We differentiate between the management part in which the order parameters are further defined (trading strategy selection, order timing and size, venue selection) and order processing part which manages the interaction with trading venues.
- **Trading venues**: provide the venues on which trade execution takes place;
- **Data analytics**: their role is to support both pre-trade and post-trade analytics. We differentiate between the high frequency part which is tightly coupled with data originating from markets and the low frequency part which feeds both upstream and downstream activities related to determining or refining order parameters. Analytics can also include data from other sources (such as news, company data etc.)

Even such a representation is still very simplified as there could be many buy-sides using multiple sell-side brokers which themselves are operating on multiple markets resulting in highly complex and interconnected ecosystem (hence the term "swarm" trading is increasingly being used).

6.2 Research Gaps

The study showed there is a need for new research work to support recent industrial innovations particularly in the following areas:

1. Important changes in the regulatory frameworks (particularly the MiFID regulations in Europe) has resulted in a proliferation of software services and component technologies that can perform various functions in the trading cycle. Therefore, it is challenging to design new machine learning models that will perform a very narrow function the trading cycle without a good understanding of the implications on other parts of the system. For example, it is hard to assess whether such models will violate compliance requirements with relevant regulations
2. The rise in the availability, variety and volume of data associated with the different models used during the trading cycle, as well as a diversification in the types of instruments that can be traded, combined with the increasing use of sophisticated machine learning techniques (particularly deep learning) in pre-trade and post-trade analytics has resulted in formidable challenges when it comes to implement faster, more "understandable" and more accurate decision-making.
3. As most existing models and systems have been designed to operate with static data, there is now a need for new cost-effective solutions that are able to acquire, process and analyse real-time data. In addition, the need to integrate several systems (possibly from different vendors) has revealed many issues related to workflow design and efficiency. Finally, the tendency to use

machine learning techniques to assist in selecting the most appropriate trading strategy based on a user's parameters is set to continue.

4. The risks associated with automated trading are poorly understood especially the quantification of risk which requires a new risk management culture to be established in financial institutions and regulators.

Figure 3. Information Flow in an Electronic Trading System

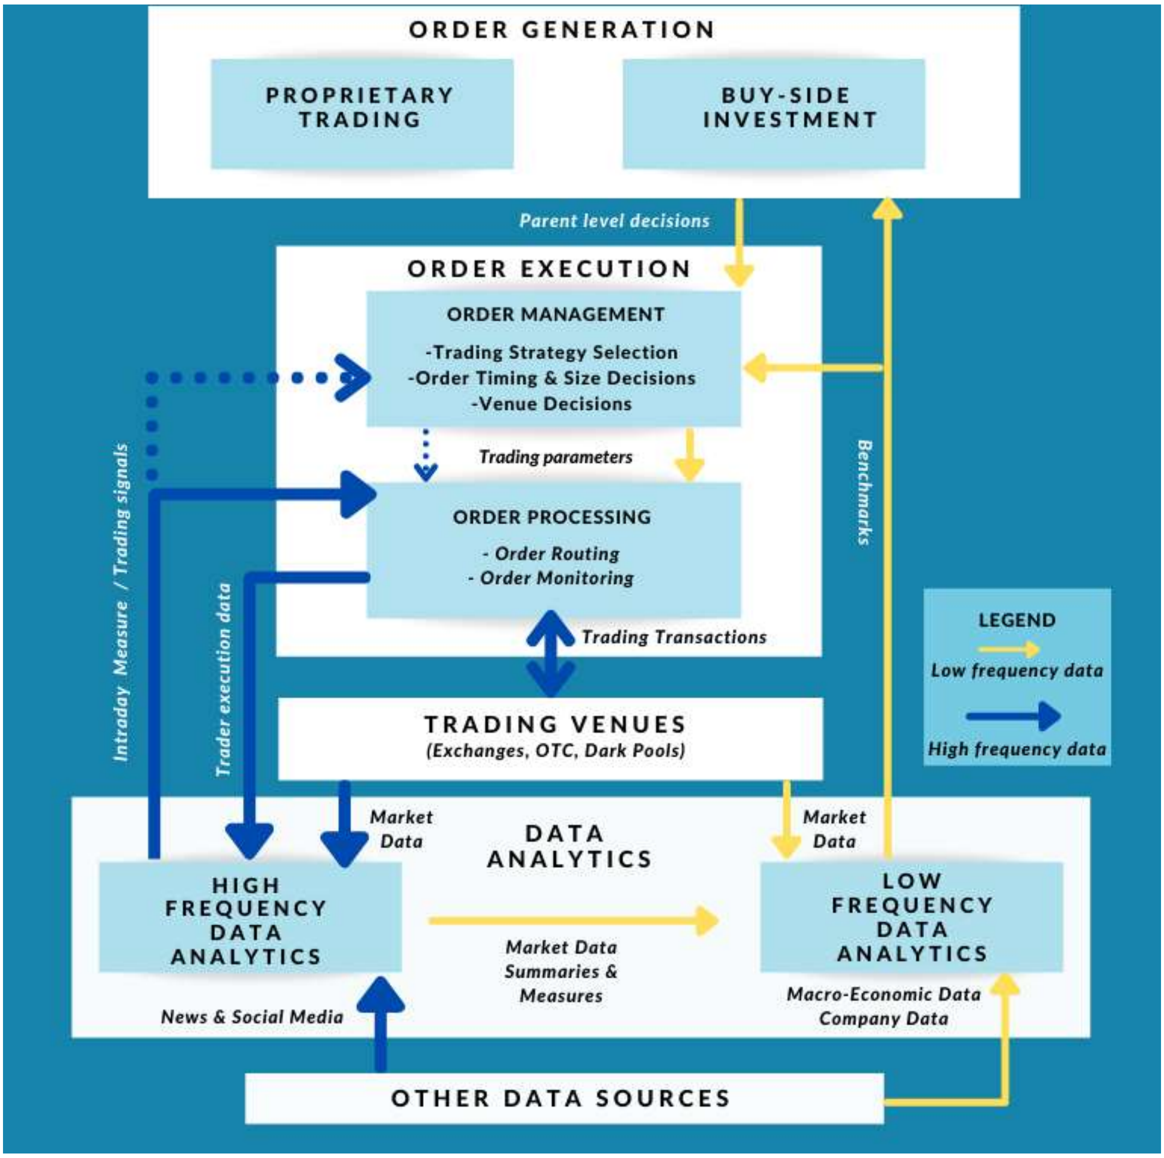

6.3 Future research areas and actions

For the research community, we identify the following important research topics that are defined around the challenges outlined earlier:

Information representation for machine-learning powered analytics: Feeding hundreds of low-level indicators into machine learning algorithms is impractical and empirical results have confirmed a severe drop in the performance of trading systems in such cases. New research into better ways of representing and abstracting data and particularly textual data is needed. New types of information that can realistically represent the richness and diversity of existing institutional environments, regulations and trading behaviours need to be defined, including behavioural data e.g., one that reflects traders' emotions. This raises many challenges such as ensuring data quality issues, managing provenance information needed for explainability as well as organizing metadata when combining data from multiple sources.

Developing more sophisticated machine learning techniques: Existing models will need to be adapted to include constraints, incentives and biases of decision makers (fund managers, traders, etc) and their impact on market valuations. In addition, new models for quantifying risk are urgently required. For these reasons, interest in the use of deep learning methods is growing, for their ability to detect and exploit interactions in the data that are not captured by any existing financial economic theory. In addition, the use of machine learning methods as a part of a bigger system in which other models co-exist has many practical advantages. We expect to see more work on how to engineer complex machine learning systems within an organisation than on designing new machine learning techniques or customizing existing ones.

Automatic trading strategy elaboration and execution: The use of machine learning techniques in the elaboration of a trading strategy and its various parameters such as order size and venue choice, taking into account a wide range of goals, time horizons and benchmarks is still in its infancy. Trying to capture individual/specific effects for a large number of financial instruments rises the number of predictors and the level of noise which reduces the performance of the trading system. Deeper empirical studies involving the hyper-optimization of all parameters are necessary. It is expected that optimization of trading processes will also be increasingly automated, guided by applying machine learning techniques to learn from experiences captured in vast amounts of execution data.

For organisations, we recommend the following actions:

- Recognise the need to have a proper governance and implementation framework to guide the implementation of machine learning in your organisation and work towards developing industry standards and guidelines in this area with similar organisations.
- Invest in explainable solutions, and by ensuring the risk management boards have inter-professional membership with stakeholders from the whole organisation including IT support.
- Support multidisciplinary educational and research opportunities across different fields such as computer science, finance and econometrics that can facilitate the development of expertise and education resources in this area.

7 Acknowledgements

We wish to acknowledge Plato Consulting and in particular Mike Bellaro for sponsoring this project. We also would like to thank Professor Carole Comerton-Forde, Professor Peter Gomber, Professor Brahim Saadouni, Dr Giuseppe Nuti, Dr Ali Behnaz, Rochak Lamba and Dr Kingsley Jones for their help and advice related to this project.

8 References

[1] Bifet, A., Gavaldà, R., Holmes, G., and Pfahringer, B. (2018). Machine Learning for Data Streams, with Practical Examples in MOA, MIT Press, DOI: https://doi.org/10.7551/mitpress/10654.001.0001

[2] BoE (2019). “Machine learning in UK financial services.” [Online]. Available: http://www.bankofengland.co.uk/report/2019/machine-learning-in-uk-financial-services.

[3] Cliff D. and Treleaven, P. (2010). Technology trends in the financial markets: A 2020 vision, UK Government Office for Science’s Foresight Driver Review on The Future of Computer Trading in Financial Markets – DR 3, October 2010.

[4] Das, S. R. (2019) “The future of fintech,” FINANCIAL MANAGEMENT, vol. 48, no. 4, pp. 981–1007, Dec. 2019, doi: 10.1111/fima.12297.

[5] Fischer, T. & Krauss, C. (2017). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research. 270. 10.1016/j.ejor.2017.11.054.

[6] Gomber, P. and Zimmermann, K. (2018). “Algorithmic Trading in Practice,” The Oxford Handbook of Computational Economics and Finance, Feb. 2018, doi: 10.1093/oxfordhb/9780199844371.013.12.

[7] González, A., Guldrís-Iglesias, F., Colomo-Palacios, R., Gómez Berbis, J., Jiménez-Domingo, E., Alor-Hernández, G., Posada-Gómez, R. & Robles, G. (2010). Improving Trading Systems Using the RSI Financial Indicator and Neural Networks. 6232. 27-37. 10.1007/978-3-642-15037-1_3.
[8] Huck, N. (2019). "Large data sets and machine learning: Applications to statistical arbitrage," European Journal of Operational Research, vol. 278, no. 1, pp. 330–342, Oct. 2019, doi: 10.1016/j.ejor.2019.04.013.
[9] JPMorgan (2020). JPMorgan e-Trading 2020 survey: https://www.jpmorgan.com/global/markets/e-trading-2020
[10] Paiva, F., Cardoso, R., Hanaoka, G., Duarte, W. (2019), Decision-making for financial trading: A fusion approach of machine learning and portfolio selection, Expert Systems with Applications, 2019 vol: 115 pp: 635-655.
[11] Preda, A. (2017). Noise: Living and Trading in Electronic Finance, 1 edition. Chicago: University of Chicago Press, 2017.
[12] Rabhi, F. A., Mehandjiev, N., & Baghdadi, A. (2020, August). State-of-the-Art in Applying Machine Learning to Electronic Trading. In International Workshop on Enterprise Applications, Markets and Services in the Finance Industry (pp. 3-20). Springer, Cham.
[13] Weber, J. and Riera, M. (2019). Welcome to the New World of Equity Trade Execution: MiFID II, Algo Wheels and AI. Targeted News Service, Greenwich Associates News Release, 23 April 2019.
[14] Weng, B., Lu, L., Wang, X. & Megahed, F. & Martinez, W. (2018). Predicting Short-Term Stock Prices using Ensemble Methods and Online Data Sources. Expert Systems with Applications. 112. 10.1016/j.eswa.2018.06.016.