Machine Learning in Algorithmic Trading

Application by Dutch Proprietary Trading Firms and Possible Risks

Last updated at 28-09-2023

Continue reading

AFM

Table of contents

Executive Summary 3

01 Introduction 4
- Aim study 5
- Approach study 5
- Scope or limitations study 5
- Structure of this report 5

02 Terminology matters 6
- 2.1. Algorithmic trading and high-frequency trading 6
- 2.2. Machine Learning in algorithmic trading 7
- 2.3. Machine learning based trading algorithm 9

03 Observations 12

04 Risks 14
- 4.1. Explainability 14
- 4.2. Market Manipulation 15

05 Discussion points 17

Annex 1 Definitions 18

Annex 2 Risks of Machine Learning in Trading 19

Annex 3 RTS 6 articles possibly affected by lack of explainability 23

Executive Summary

Parallel to the success of programs such as Deep Blue and AlphaGo, developments in artificial intelligence and machine learning grabbed the attention of industries worldwide, and therefore the attention of supervisors and policy makers. Since then, think-thanks, academics and supervisors have written extensively about machine learning, and its implications for the financial markets.

However, to the best of our knowledge, few studies have been published about the *actual* use of artificial intelligence or machine learning in algorithmic trading. Also, few supervisors have shared concerns about risks relevant to *conduct* supervisors, as opposed to risks for financial markets in a more general sense.

This AFM aims to do exactly that with this publication: report about the actual use of machine learning as reported by a subset of Dutch proprietary trading firms, and report about the possible risks relevant to its supervision. The aim of this publication is to contribute to the public debate, inform academia and other supervisors. In addition, the AFM uses the findings in this study to focus its supervision on the most relevant risks to its supervision.

Please note that observations in this study are based on surveys sent to – and interviews held with – a subset of Dutch propriety trading firms. All trading firms used algorithmic trading. Note that the observations in this study are not necessarily representative of the use of machine learning in algorithmic trading by other segments in the financial markets (e.g., brokers executing orders for clients). Also, the findings represent a snapshot of the state of the market in the year 2022.

The AFM observes that (see section 2 and 3):

1. Clear terminology is required to account for the many nuances in algorithmic trading
2. Machine learning is applied on a large scale in algorithmic trading
3. Many machine learning models used in algorithmic trading try to predict the price of a financial instrument
4. Machine learning models look primarily at order book data, no fundamental information, and use 100-1.000 features
5. Trading firms heavily rely on supervised learning, not (yet) reinforcement learning
6. Trading firms find explainability of models less important than performance
7. Trading firms see risks of reinforcement learning based trading algorithms to learn unintentional and negative trading behaviour

Also, the AFM observes several possible risks from the use of machine learning in algorithmic trading (see section 4). Two risks that are especially relevant to a conduct supervisor such as the AFM are:

1. **Lack of explainability of machine learning models** poses challenges for trading firms to comply with organisational requirements regarding algorithmic trading
2. **Increased risk of market manipulation.** Reinforcement learning could increase the risk of trading algorithms learning unintended, negative trading behaviour, while the implicit use of machine learning could make trading algorithms more susceptible to falling pretty to manipulation

01 Introduction

It was May 11, 1997. World chess champion Garry Kasparov, to this day considered one of the greatest chess players of all time, conceding the last game in a match to his opponent: IBM supercomputer Deep Blue. Man competed against the machine. The machine won.

Deep Blue was a chess-playing 'expert system': a computer system emulating the decision-making ability of a human expert. Expert systems are designed to solve complex problems by reasoning through bodies of knowledge, represented mainly as a body of if–then rules.

Deep Blue's rules were designed by human experts in chess. In this sense, the program acted like a human expert would. The difference between Deep Blue and a human expert, however, was that Deep Blue could evaluate many more positions than any human ever could: 200 million positions per second. And this paid off.

Almost twenty years later the computer program AlphaGo came to the stage. AlphaGo then competed against legendary Go player Mr Lee Sedol. AlphaGo's 4-1 victory in Seoul, caught headlines worldwide.

Go is a profoundly complex game, much more complex than chess. There are an astonishing 10 to the power of 170 possible board configurations – more than the number of atoms in the known universe. Hence it was impossible to capture each board configuration in a body of if-then rules, as had been tried with Deep Blue in the game of chess. A different approach was required.

AlphaGo's developers used so-called 'deep neural networks' and 'reinforcement learning', both types of 'machine learning'. These models take a description of the Go board as an input and process it through several network layers containing millions of neuron-like connections. Then the developers had it play against different versions of itself thousands of times, each time learning from its mistakes.

Over time, AlphaGo improved, became increasingly strong and better at learning and decision-making.

Not only did AlphaGo win, but it also invented some winning moves, several of which were so surprising that they upended hundreds of years of wisdom. The reason AlphaGo could surprise human experts was precisely because it didn't follow strict human instructions, as Deep Blue did, but learned from its own experience playing the game.

Parallel to the success of programs such as Deep Blue and AlphaGo, developments in artificial intelligence (AI) and machine learning (ML) grabbed the attention of industries worldwide, and therefore the attention of supervisors and policy makers. Since then, think-thanks, academics and supervisors have written extensively about machine learning, and its implications for the financial markets. See for example the OECD's report on opportunities, challenges, and implications for policy makers because of the application of artificial intelligence in finance.[1]

1 See oecd.org/finance/financial-markets/Artificial-intelligence-machine-learning-big-data-in-finance.pdf

Aim study

However, to the best of our knowledge, few studies have been published about the *actual* use of artificial intelligence or machine learning in algorithmic trading. Also, few supervisors have shared concerns about risks relevant to *conduct* supervisors, as opposed to risks for financial markets in a more general sense.

The AFM aims to do exactly that with this publication: report about the actual use of machine learning as reported by a subset of Dutch proprietary traders, and report about the possible risks relevant to its supervision. The aim of this publication is to contribute to the public debate, inform academia and other supervisors. In addition, the AFM uses the findings in this study to focus its supervision on the most relevant risks to its supervision.

Approach study

The AFM started this study by reading literature on the topic and coming up with a shortlist of possible risks relevant to the supervision of a conduct supervisor on the financial markets. Next, we designed a detailed survey using the input of several trading firms. The objective of the survey was to get quantitative information from some Dutch proprietary trading firms about their use of machine learning in algorithmic trading, and to report on the risks they see.

The surveys were filled out by several Dutch algorithmic trading firms, and subsequently discussed in interviews with subject matter experts from the firms.

Scope or limitations study

The reader should bear in mind that:

- The observations in this study are based on information provided to the AFM by several large, Dutch, proprietary trading firms that use (almost exclusively) algorithmic trading.
- The observations are not necessarily representative for the algorithmic trading industry in general, especially not for algorithmic trading by firms not trading on their own account (e.g., brokers that execute orders for clients).
- The observations and risks relate to the *inner workings* of trading algorithms, not necessarily to the speed at which they trade.
- The observations are based on information provided to the AFM by trading firms. The AFM is not able to independently verify all claims.
- Technology is subject to constant development. Therefore, the observations today (2022) might differ from observations in the future. The same applies to priorities in risk assessment.

Structure of this report

Section 2 explains some concepts that are relevant to understand the application of machine learning in algorithmic trading. We pay special attention to the difference between a ‘machine learning algorithm’ and a ‘trading algorithm’ and show how the two meet in a ‘machine learning based trading algorithm’. Section 3 lists the AFM’s main observations about the use of machine learning in algorithmic trading. We also note trading firms’ assessment on potential risks. Section 4 focuses on possible risks from the use machine learning in algorithmic trading, focusing on two risks in particular. Lastly, the AFM raises a few points that it considers relevant to both regulators and market participants, and hope they might stimulate discussion.

02 **Terminology matters**

In this section we explain some notions that are important to understand the application of machine learning in algorithmic trading. First, we explain algorithmic trading and high-frequency trading, and how the two differ. Then we give a brief overview of machine learning techniques.

In our discussion with supervisors and trading firms, the AFM observed that the terms 'machine learning algorithm' and 'trading algorithm' tend to be used interchangeably. Furthermore, the AFM observes that supervisors tend to think mostly of 'reinforcement learning' when thinking about machine learning in the context of algorithmic trading.

In section 2.3. we aim to show how a machine learning algorithm is different from a trading algorithm, but that the former can be an intrinsic part of the latter. We point to the 'implicit' and 'explicit' use of machine learning in trading algorithms. Our hope is that, by pointing out the more implicit use of machine learning in algorithmic trading, the reader will see that machine learning has implications for the financial markets that go far beyond reinforcement learning only.

2.1. Algorithmic trading and high-frequency trading

The algorithmic trading universe is vast, and algorithms come in many shapes and sizes. Trading algorithms can use different techniques 'under the hood', ranging from bodies of if-then rules to advanced artificial intelligence. They are used by many different actors, ranging from pension funds to proprietary traders. Some trading algorithms are very fast and need to be so, while for others speed is less important.

When thinking and talking about algorithmic trading, terminology matters a great deal. There is a need to speak the same language and account for the many nuances in algorithmic trading, especially when it comes to analyzing the impact and risks of algorithmic trading for the financial markets.

We start by defining three terms:
Execution algorithms are algorithms that aim to execute an order. The decision to invest in the financial instrument is taken elsewhere. These algorithms are often used to place large orders in the market as to minimize price impact. An example would be a Volume-Weighed Average Price (VWAP) algorithm.

Trading algorithms or investment decision algorithms aim to automate a strategy, and automatic execution is part of that. In contrast to execution algorithms, a trading algorithm does take the investment decision. The algorithm can be programmed via a body of if-then rules to initiate an action if certain conditions are met, but it can also use a machine learning model to detect trading opportunities (as we will see below). An example would be an algorithm that seeks to optimize a portfolio's exposure and automatically execute the strategy.

Algorithmic trading

Automated *Execution*	Automated *Trading*
Goal: execute an order (intelligently)	**Goal: automate a trading strategy**
Hedge funds, pension funds, investment banks, brokers	Hedge funds, trading desks of investment banks, proprietary trading firms.
Investment decision is taken elswhere	Investment decision is taken by algorithm

Another concept associated with algorithmic trading, and often (wrongly) used as equivalent to algorithmic trading, is **high-frequency** trading (HFT). With HFT, a trading system analyses market data at a very high speed, and sends large numbers of orders or revises these orders within a very short time span in reaction to this analysis.

HFT is not a strategy, it is a technology with which trading strategies are executed.[2] So some algorithmic trading firms use HFT, while others don't.

2.2. Machine Learning in algorithmic trading

In this section we introduce the notion of machine learning in the context of algorithmic trading. The aim is to give the reader a sense of the different types of machine learning and how they can be used in algorithmic trading.

Introduction

Broadly defined, we take machine learning to consist of sets of rules that use data to automatically 'get better' at performing a particular task.[3]

With 'getting better' in the context of algorithmic trading, one could think of getting better in predicting or estimating a share price ('Is the share price going up in the next second?', 'What is the theoretical price of an option?'), which is the domain of supervised learning, or to make more profit or to achieve better execution, which is the domain of reinforcement learning (more on this later).

The machine learning algorithm does so by 'learning'. By learning in this context, the AFM does not refer to the mental process of learning as humans do. Learning in the context of machine learning refers purely to the process of getting better.

2 Algorithmic trading (afm.nl)
3 One could argue that this definition does not quite fit unsupervised learning. After all: while supervised learning and reinforcement learning might get better at predicting a future price or picking the best action in any state, unsupervised learning has a less clear benchmark to measure improvement. Hence: one could add 'or finds patterns in data'.

'**Features**' play an important role in the context of machine learning, and this term will be used frequently within this study. A feature is any measurable or quantifiable characteristic that represents some relevant phenomenon in the context of the modelling problem and is used as an input to a machine learning model. Think of a feature as a numerical representation of (relevant parts of) the state of the world which can be processed by a machine learning model. A key component of the machine learning process is to determine the right set of features for the task at hand.

Example:
Trading firm X wants to predict the price of share XYZ 1 second from now. The firm believes that the current volume in the order book and the price trend over the past 10 seconds contain some valuable information about where the price will be heading in the next second. Hence trading firm X chooses 'volume in order book' and 'price trend over past 10 seconds' as features in its prediction model.

Machine Learning techniques

The process a machine learning algorithm uses to learn and the type of output it obtains depend on the type of technique applied. The field of machine learning is often divided into three types of techniques:
1. Unsupervised learning
2. Supervised learning
3. Reinforcement learning

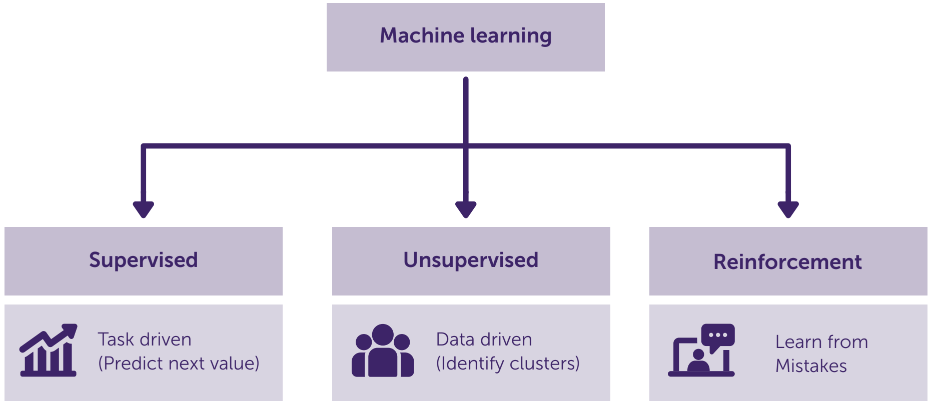

Source: [1. Machine Learning in Finance: The Landscape - Machine Learning and Data Science Blueprints for Finance [Book] (oreilly.com)]

Based on conversations with academics and Dutch algorithmic trading firms, the AFM observes that supervised learning is the current 'go-to' technique for machine learning in the Dutch proprietary algorithmic trading industry, but that reinforcement learning might be the 'future' (see section 3).

We will briefly discuss these three different types of techniques in the context of algorithmic trading.

Unsupervised learning

An unsupervised learning algorithm is a type of algorithm that learn patterns in data. Common applications in the trading industry are clustering ('What groups of financial instruments tend to have similar properties?') or dimensionality reduction ('Can we reduce 1.000 features to 100 features, while still predicting the share price reasonably well?'). These techniques are mainly used in the exploration stage of new trading strategies.

Example:
Trading firm X wants to know which shares tend to behave similarly. That is: they would like to divide shares into groups A, B, C, D, so that shares which belong to the same group have similar levels of volatility, trading volume or other characteristics. The clustering can in turn be used to assign different risk profiles to the different groups of shares. Trading firm X uses a k-means clustering algorithm to do so.

Supervised learning

A supervised learning algorithm learns by repeatedly adjusting the weights[4] of its features so that its output deviates less and less from the output that the algorithm is supposed to give. To know what output the algorithm is supposed to give (the 'labels'), the algorithm is handed many examples. That is: the algorithm is shown what the output should be for a given set of features.

Now, via gradually changing its weights, the supervised learning algorithm is supposed to learn the relationship between the features and the output, and eventually it should be able to infer the output (or predict what the output will be in the future) by only looking at the features. Once that point is reached, a trading firm can use the algorithm to improve its trading performance.

A common application of supervised learning by proprietary algorithmic trading firms is to predict the price of a financial instrument (see section 3).

Example:
Trading firm X wants to predict the midpoint price of share XYZ 1 second from now. The firm thinks that some features might have predictive value, e.g., the historical trend in the share price might tell something about where the price will be heading. They apply a linear regression model to many rows of historical data and obtain a model which has learned the relation between trend in share price (and many other features) and future share price. They use the predictions of this 'trained' model in their trading algorithms.

4 Weight can be thought of as the mathematical equivalent to the contribution of a feature to the output of the model.

Reinforcement Learning

In contrast to supervised learning, a reinforcement learning algorithm does not learn from labelled data, but via **trial and error**. That is: one implements a reinforcement learning based trading algorithm in a simulation environment (or in a real market for that matter), and the algorithm is supposed to learn over time to choose the optimal action in each state of the market.

The algorithm does so by measuring how much each action in each state has contributed to some overarching objective, e.g., profit or execution costs. If an action did well (i.e., cancelling an order is followed by high profits at the end of the trading session), then the trading algorithm is more likely to pick the action the next time it faces the same state of the market. Was the effect negative (e.g., sending in a large market buy order is followed by high execution costs at the end of the trading session), then the action will be picked less frequently in that state of the market.

So, reinforcement learning is not concerned with predicting a future share price, for example, but with choosing the actions the trading algorithm should take in order to achieve some objective. It does not do so by comparing its action to some set of actions the algorithm is 'supposed' to take (such as the predictions an algorithm is 'supposed' to make in the case of supervised learning), but by finding out over time and through experience what the optimal action is.

The result is a model that chooses the best possible action (i.e., the action that historically has done best) in every state of the market.

Example:
Trading firm X wants to optimise profits. Being a market maker, it would like to know at what price level to send in a new buy order, given the state of the order book at the current point in time. It can pick actions such as: bid level 1, bid level 2, ..., bid level 100, etc. They trained a model using reinforcement learning. Given the current state of the order book, the model has learned that sending in a buy order at the 5th-best level in the order book has resulted in the most profitable runs in the past. Hence, the trading algorithm will send in a buy order at the 5th price level.

2.3. Machine learning based trading algorithm

In our conversations with trading firms the AFM observed that it is very important to use the right terminology when talking about machine learning in the context of algorithmic trading. For example: asking a trading firm if it uses machine learning in its trading, the firm might say 'No'. If we would ask the same firm if its trading algorithms in some way use the output of a machine learning technique to determine some trading parameters such as 'price' or 'volatility' in some model (which we might equate to 'using machine learning in trading algorithms'), the answer might be 'Yes'.

The main point of confusion seems due to the fact that both a trading algorithm and a machine learning algorithm are instances of an 'algorithm', yet have very different functions. A machine learning algorithm doesn't do any trading, and a trading algorithm doesn't do any machine learning. A trading algorithm might use the *output* of a machine learning algorithm to determine one or more components in the trading algorithm, but the trading algorithm itself does not do any learning.

Also, the AFM observes that most supervisors, and some trading firms, tend to think solely of reinforcement learning when talking about machine learning being used in trading algorithms.

The AFM aims to show the wider implications and impact of machine learning in algorithmic trading by pointing out also the more implicit use of machine learning in trading algorithms.

We will define the term '**machine learning <u>based</u> trading algorithm**' to refer to any trading algorithm that uses (the output of) a machine learning model. A machine learning based trading algorithm can use machine learning either *implicitly* or *explicitly*.

Implicit

The trading logic of a machine learning based trading algorithm can be very similar to the logic of a 'traditional' or non-machine learning based trading algorithm. The only difference might be that certain numbers that had been pre-defined or hard-coded in the trading algorithm (see example below), are now obtained via a machine learning model. So, a machine learning component is added to a traditional trading algorithm, while not fundamentally changing the trading algorithm's decision logic.

In such cases, we call the trading algorithm implicitly machine learning based.

We present a very simple trading algorithm that sends in a buy order if it assumes the price of the share 1 second from now will be higher than it is now. Note that this is a trading algorithm, not a machine learning algorithm:

Trading algorithm

```
If price_1_second_from_now_higher(x) then:
      send_buy_order(x)
End if
```

Prediction (obtained via machine learning or not)

At this point in time, we don't have sufficient information to determine if this is a machine learning based trading algorithm or not. In green we have the component that could turn this trading algorithm into a machine learning based trading algorithm. This all depends on how the price_1_second_from_now_higher(x) function arrives at its prediction.

Option A: Expert-based

Suppose a group of experienced traders decide that the following are reasonably good predictors of the share price 1 second from now:

- Trend over past 10 seconds
- Number of liquidity taking buy transactions minus liquidity taking sell transactions in past y seconds
- Trend in related instruments

Furthermore, they define for each predictor certain thresholds, meaning: if the value would go above/below a pre-determined number, then the traders think the price 1 second from now will be higher. For example: If the price has increased more than 1% over the last 10 seconds, then the traders assume the price will keep going up in the next second. They formalise this into a body of if-then rules, which basically says that if the price went up more than 1% over the last 10 seconds, then we predict the price will go 'Up' in the next second:

```
Def price_10_second_from_now_higher(x):
      If trend_past_10_seconds > 1% or ….
            Return "True"
      Else:
            Return "False"
      End if
```

We plug this function into our trading algorithm and (assuming our traders have indeed isolated good predictors) the predictions will be reasonably accurate.

Option B: Machine learning based

The trading firm that uses this trading algorithm thinks that a machine learning model might give them better price predictions, and in turn a more profitable trading algorithm.

They have trained a logistic regression model that uses 1.000 features to predict if the price 1 second from now will be higher:

```
Def price_1_second_from_now_higher(x):
      If logistic_regression_model(x) == 1^6 then:
            Return "True"
      Else:
            Return "False"
      End if
```

The trading firm plugs this trained model into its trading algorithm.

We see that the decision logic of the trading algorithm remains the same. It still decides to send in a buy order if it predicts the price to be higher 1 second from now. The way it arrived at the prediction, however, is different: in case of Option B the output of a machine learning model determines whether to send an order or not, while in case of Option A the output of a non-machine learning model (read: a body of if-then rules) does. Therefore, we call the trading algorithm using option B machine learning based, and option A not.

The AFM observes that proprietary trading firms in scope of our study make frequent use of such implicitly machine learning based trading algorithms (see section 3).

Explicit

In the case of an implicitly machine learning based trading algorithm as defined in the previous section, there is a relatively clear distinction between the machine learning algorithm and the trading algorithm. One produces an output that is used by the other.

In the case of an explicitly machine learning based trading algorithm, this distinction is much less clear. By an explicitly machine learning based trading algorithm the AFM refers to any trading algorithm that uses (the output of) a machine learning model, where the latter is not restricted to a well-defined stage of its set of instructions. One could say that the decision logic ('What action to take in which state?') of the trading algorithm is (largely) determined by the machine learning model. The prime example of an explicitly machine learning based trading algorithm is a reinforcement learning based trading algorithm, that has learned which action to take in any state of the market.

Compare the following example to the implicitly machine learning based trading algorithm above. In the previous example the programmer explicitly prescribed which action to take ('send_buy_order') if a condition was met. In this example, the reinforcement learning algorithm decides which action to take if a condition (read: a 'state of the market') is met. This is a fundamental difference: while the programmer hands the reinforcement learning algorithm a couple of actions to choose from, *it does not know up front* which action the model and therefore the trading algorithm will take.

Trading algorithm

```
current_state = get_current_state_of_market()
take_best_action_in_state(current_state)
```

Optimal action (obtained via Reinforcement Learning)

It is the latter use of machine learning in algorithmic trading that supervisors often think of. Yet, as we will see in section 3, this is not (yet) the main application of machine learning in trading algorithms.

03 Observations

In this section we describe the AFM's main observations about the use of machine learning in algorithmic trading.

Recall that these observations are based on surveys sent to – and conversations had with – some Dutch proprietary trading firms. Also, recall that the observations are based on information provided to the AFM by trading firms, and are not necessarily representative of other segments of the algorithmic trading industry (e.g., brokers executing orders for clients). See section 1 for the scope of our study.

1. Machine learning is applied on a large scale in algorithmic trading

Trading firms tell the AFM that machine learning is implicitly or explicitly used in 80%-100% of their trading algorithms. This percentage is higher than the AFM expected, and shows that machine learning is no hype, but an intrinsic part of the business models of many trading firms nowadays.

2. Many machine learning models used in algorithmic trading try to predict the price of a financial instrument

Price prediction is a clear example of how machine learning can be implicitly used in trading algorithms – see section 2.3.

Firms tell the AFM that machine learning is most heavily used in liquid asset classes, such as equity and futures. Machine learning is also used in other asset classes (e.g., options), although its application tends to be different than its application in equity and futures. Where machine learning models in liquid asset classes are mostly used to predict the future price of a financial instrument, machine learning in less liquid asset classes is more often used to calibrate parameters such as volatility, which are in turn used as parameters in pricing models of trading algorithms.

3. Machine learning models look primarily at order book data, no fundamental information, and use 100-1.000 features

Trading firms tell the AFM that:

- Their machine learning based trading algorithms look at features such as:
  - o Volume in order book
  - o Price trend
  - o Volatility
  - o Order book imbalance
- Models use many permutations of the same features (e.g., volume in order book on the first level, volume in orderbook on the second level, etc.).
- Models used in one financial instrument use features computed on other financial instruments, asset classes and trading venues (e.g., order book data of a share on CBOE is important to predict the price of the same share on Euronext).
- Models use in the range of 100-1.000 features.
- Features computed over the recent past (sub second) tend to be more predictive than features computed over longer time horizons (e.g, minutes).
- Models do not look at fundamental information (e.g., yearly reports).

4. Trading firms heavily rely on supervised learning, not (yet) reinforcement learning

The trading firms in scope of this study tell the AFM that:

- Most machine learning models in use are supervised learning models.
- Unsupervised learning models are used, but mostly in the pre-trading phase (e.g., feature selection).
- Reinforcement learning based trading algorithms are not used in practice yet. However, most agree that reinforcement learning might be a promising technique in the future. Recall that reinforcement learning allows trading algorithms to learn through trial and error which action to take in any state of the market.
- Trading firms tend to prefer simpler supervised learning models such as linear and logistic regression to more complex models such as artificial neural networks.
- The actions that a machine learning based trading algorithm can take are hard-coded and limited, e.g., send a buy order or sell order, or cancel an existing order. So, most trading algorithms are implicitly machine learning based, as defined in the previous section.
- Machine learning models are retrained periodically, ranging from weeks to months. Retraining a model could potentially change the way a trading algorithm responds to market conditions.

5. Trading firms find explainability of models less important than performance

The trading firms in scope of this study tell the AFM that they are not much concerned about a potential lack of explainability of their machine learned based trading algorithms. Firms stress that performance or predictably of trading algorithms is more important to them than explainability of the models. Furthermore, firms tell the AFM that any trading algorithm, machine learning based or otherwise, should be judged on its output (i.e., orders, transactions, deletions, etc.), not how that output came about.

So irrespective of how a trading algorithm arrives at its trading decision, most trading firms state it might be sufficient to monitor the *conduct* of the trading algorithm, not necessarily its *inner workings*. It is in this sense, firms claim, that machine learning based trading algorithms do not differ from non-machine learning based trading algorithms.

See section 4 for the AFM's views on the risks of lack of explainability of machine learning based trading algorithms.

6. Trading firms see risks of reinforcement learning based trading algorithms to learn unintentional and negative trading behaviour

The trading firms in scope of this study tell the AFM that they do see limited risks of a supervised learning based trading algorithm (e.g, the price prediction example in section 2) to learn unintended and negative trading behavior, or to manipulate the market. However, trading firms do see risks of reinforcement learning based trading algorithms to learn unintended and negative behavior. Trading firms tell the AFM that this is one of the reasons they are at present hesitant to apply such techniques in practice.

See section 4 for the AFM's views on the risks on market manipulation of machine learning based trading algorithms.

04 **Risks**

In this section we describe the main risks of the use of machine learning in algorithmic trading, from the perspective of a conduct supervisor.

The AFM identifies several potential risks of the use of machine learning in algorithmic trading to the fair and orderly functioning of financial markets. See Annex 2 for some more potential risks of the use of machine learning in algorithmic trading to the integrity of the financial markets. Based on an analysis of our supervisory tasks and relevant legal articles, the AFM considers two risks especially important: lack of explainability of machine learning based trading algorithms, and risks of market manipulation of machine learning based trading algorithms.

Recall that the AFM is the conduct supervisor on the Dutch financial markets. As such, risks relevant to its supervision might differ from risks relevant to a prudential supervisor, such as the Dutch Central Bank.

4.1. Explainability

As pointed out in section 3, trading firms in scope of this study tell the AFM they are not much concerned with any lack of explainability of trading models. However, the AFM does notice risks due to a lack of explainability of machine learning models used in algorithmic trading.

The AFM acknowledges that more complex machine learning models can outperform simpler models (e.g., they might be better in predicting a future share price). However, more complex models might also be less easy to interpret, explain and control by trading firms. Similarly, adding more features to a model might improve its performance, but might also make it more difficult to understand or to explain the effect of a feature on the algorithm's trading behavior.

This added complexity has direct ramifications for trading firms – see Annex 3 for a selection of relevant RTS6-requirements that might be impacted by lack of explainability. For one, if a trading algorithm is suspected to cause disorderly trading conditions, the trading firm should be able to explain to the supervisor how the algorithm's trading decision came about. This information is important for a supervisor, since it allows the supervisor to assess if the resulting behavior was intended or unintended. Note that, in a more general sense, a trading algorithm should not act in an 'unintended manner' (RTS 6 Article 5(4) – see Annex 3).

Furthermore, the risk and compliance staff of a trading firm should have sufficient knowledge of the algorithmic trading strategies of the trading firm, and sufficient authority to challenge trading staff. The more complex trading models become, the more difficult it will be for risk and compliance staff to challenge developers on the implications of using specific features or machine learning models (Annex 3 RTS 6 Article 3(4)).

Also, the more features a model uses and the more complex a model becomes, the more difficult is for any human to deduce how the trading algorithm will behave under any given circumstance a priori. The AFM already sees that it is difficult for some trading firms to be sure how their machine learning based trading algorithm will behave *before* deployment in a testing or a real environment. This puts an increasingly large burden on trading firms to test their trading algorithms properly, in order to be sufficiently sure the algorithm will behave as intended and not cause any disorderly trading conditions, *before* deployment on the real markets.

The latter requires a testing environment that is sufficiently realistic to account for the interaction with other trading agents (e.g., other trading algorithms). In its 2021 study 'Algorithmic Trading – Governance and Controls', the AFM concluded that many testing environments at present are not sufficiently realistic to account for the intricate, yet important, microstructure of real trading markets.

Lastly on this matter, trading firms in scope of this study tell the AFM that that they have surveillance systems in place to detect market manipulation or unintended trading behavior if it were to take place. The AFM would like to stress that after-the-fact surveillance does not relieve firms from the responsibility to make sure their trading algorithms behave in orderly manner and as intended in the first place.

4.2. Market Manipulation

Supervisors' focus, including the AFM's, is often on detecting the agents that manipulate the market. Yet market manipulation is a function of two components: (1) agents manipulating, and (2) agents being manipulated. Hence, in case we would like to minimize market manipulation, we should also focus on making trading algorithms less receptive to manipulation.

Supervised learning

With that in mind, the AFM at present believes that implicitly machine learning based trading algorithms do not necessarily pose a greater risk to manipulate the market than traditional trading algorithms.

However, supervised learning techniques might make trading algorithms more susceptible to falling prey to manipulation. The reason for this is as follows:

We have seen that today's supervised learning based trading algorithms can use up to 1.000 features. Machine learning allows the user to extract valuable information from such features, and it is not necessary anymore to understand why a feature might have predictive value. If it does, one can add it to the model. Therefore, machine learning based trading algorithms tend to use more features than traditional algorithms, hence **more features that affect an algorithm's trading decisions.**

The effect any feature might have on the trading behavior of a machine learning based trading algorithm tends to be less transparent than in traditional trading decisions. Take a feature such as order book imbalance, which is important in the predictions of many trading algorithms.[5] This feature can itself be multiplied with many other inputs. Consequently, the effect of a change in order book imbalance on an algorithm's trading behavior becomes less obvious, allowing any negative effects (e.g., spoofing) on the trading algorithm to go unnoticed for longer. Therefore, it might become **increasingly difficult to observe when a machine learning based trading algorithm is in effect being misled or manipulated.**

We have already seen that many trading firms use a similar subset of features. Also, trading firms might use a similar dataset, hence might obtain similar information as to which features affect predictions in what manner. Knowing the importance of features, and how they affect the predictions of other trading algorithms, could allow a user **to nudge the trading behavior of other market participants into the (for the user) slightly more desirable direction.**

Combining these components (more features, the effects of features on predictions being less obvious and firms sharing a similar subset of features) could increase machine learning based trading algorithms' susceptibility to market manipulation.

Reinforcement learning

Like some trading firms (see section 3), the AFM is concerned about the implications of reinforcement learning for market manipulation. The AFM considers it likely that a reinforcement learning based trading algorithm will learn to manipulate the market, if programmed naively.

5 See this study by the AFM, the Alan Turing Institute and the Oxford Man Institute of Quantitative Finance into the predictors of the trading decisions of trading algorithms

If the objective of a machine learning based trading algorithm is to optimize some objective (e.g., profit or price impact), then – if other trading algorithms would be susceptible to manipulation and the reinforcement learning algorithm could profit from doing so – there is no reason a priori to assume the trading algorithm will not learn negative trading behavior, or even to manipulate.

The AFM notes that trading algorithms could learn to manipulate the market *even* if its developers don't want them to. Hence, good will on the side of developers is not sufficient to prevent market manipulation.

Manipulation could be restricted to a single instrument on a single trading venue, hence relatively easy to detect. Yet in principle such manipulation could be very shrewd and span multiple trading venues, asset classes or instruments, which makes it extremely difficult to detect.

The application of reinforcement learning techniques, and the implications for market supervision, is to be monitored closely. Furthermore, further research by academia or supervisors could be useful in determining the exact ramifications of reinforcement learning based trading algorithms.

05 Discussion points

The following points are related to the use of machine learning in algorithmic trading, and the AFM hopes they might stimulate discussion. The list is not meant to be exhaustive.

1. There might be substantial risks of reinforcement learning based trading algorithms to learn to manipulate the market. The AFM is told by trading firms in scope of this study that they do not yet apply reinforcement learning in practice, but we expect that these techniques are being applied in other parts of the algorithmic trading industry. The AFM believes this development is to be monitored closely. It also begs the question whether current legislation is sufficiently equipped to mitigate any risks specific to this form of AI.

2. At what point does an updated trading algorithm become a new algorithm? What if features are added or removed? What if a model is retrained? After each of these adjustments the algorithm might arrive at, for example, different price predictions, and is likely to act differently from how it acted before.

   This discussion is relevant for various reasons. One being that trading firms should – as part of their transaction reporting to the financial supervisor – report the trading algorithm used to execute a transaction. This information allows supervisors to detect trading algorithms that might behave in a wrong fashion. The same goes for testing of trading algorithms by trading firms: firms should test any new trading algorithm as to make sure they comply with RTS 6.

   Hence a clear definition of 'trading algorithm' seems appropriate.

3. One way to deal with risks of lack of explainability and market manipulation by machine learning based trading algorithms is through control frameworks (e.g., risk controls, compliance involvement, monitoring, etc.). Another might be a larger focus on (realistic) testing environments (e.g., a focus on the output of models to observe how the output might differ in different circumstances).

   Lily Bailey and Gary Gensler put this as follows: 'The supervisory focus could be shifted from documentation of the development process and the process by which the model arrives to its prediction to model behavior and outcomes, and supervisors may wish to look into more technical ways of managing risk, such as adversarial model stress testing or outcome-based metrics'

Annex 1 **Definitions**

Machine learning	*Generic term used to describe different types of techniques in which machines independently learn and aim to get better at a task or prediction. Often divided into unsupervised, supervised, and reinforcement learning*
Implicit machine learning based trading algorithm	*Trading algorithms that at some well-defined stage in its set of instructions makes use of predictions/outputs obtained via a machine learning algorithm. The logic of the trading algorithm might be very similar to a non-machine learning based trading algorithm, yet certain variables – think size and/or price or aggressiveness/ passiveness – are obtained via applying a machine learning model. An example would be a market making trading algorithm that uses (the output) of a deep neural network to determine the price of any order to send to the market.*
Explicit machine learning based trading algorithm	*Trading algorithms that use (the output of a) machine learning model, where the latter is not restricted to a well-defined stage of its set of instructions but is in fact the trading algorithm. Or equivalently: the machine-learning algorithm is not part of an overarching trading algorithm, but acts – to a large degree – autonomously in the market. An example of an explicit machine learning based trading algorithm would be a reinforcement learning based trading algorithm, that at its own discretions decides when and how to act in any particular state of the order book.*
Feature	*Any measurable or quantifiable characteristic that represents some relevant phenomenon in the context of the modeling problem and is used as an input for a machine learning model. E.g., order book imbalance when modeling price movements.*

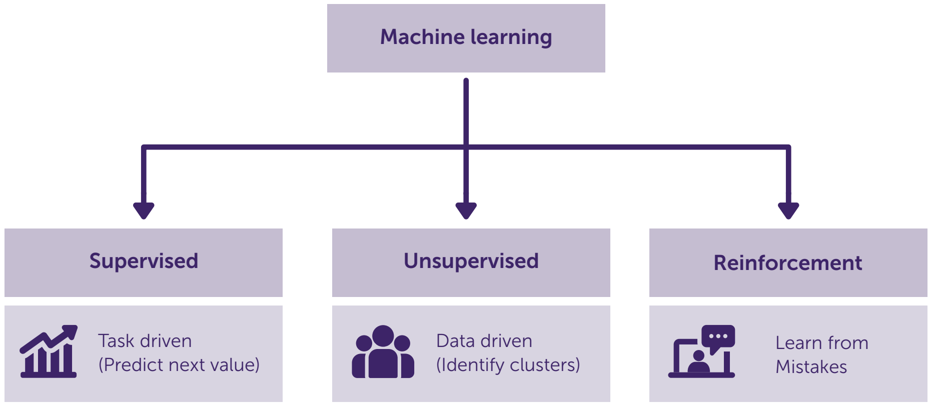

Source: [1. Machine Learning in Finance: The Landscape - Machine Learning and Data Science Blueprints for Finance [Book] (oreilly.com)]

Annex 2 **Risks of Machine Learning in Trading**

There is plenty of literature on the topic of risks of machine learning (and artificial intelligence) in the financial industry, both from an academic- and a regulatory perspective. In the literature and through conversations the AFM had with experts on the topic (both academics and market participants), the following risks recur. Please note this list is non-exhaustive.

1. **Explainability or the "black-box" nature of machine learning algorithms**
   The complexity of machine learning algorithms might make it very difficult to pin down *why* a machine learning based trading algorithm takes a certain action. This has implications for compliance with various regulatory standards.
2. **Market manipulation**
   Machine learning based trading algorithms might be more prone to (unintentionally) to manipulate the market (for example: spoofing) or colluding[6]. Also, they might be more susceptible to *being* manipulated[7].
3. **Procyclicity**
   If a proportion of the machine learning based trading algorithms respond similarly to certain market conditions, their actions might exacerbate market dynamics (think: volatility, trend, etc.), in turn causing even *more* reaction by these algorithms.
4. **Concentration risk**
   The investments required by trading firms to benefit from machine learning (in both human resources, data and technology) might make it economically feasible only to a select few, hence creating a less levelled playing field.
5. **Knowledge gap**
   Due to large investments in machine learning by especially the private sector, the difference in machine learning expertise between trading firms and supervisors (or traders and compliance personnel) might have become so large that the latter might be insufficiently able to challenge the former or detect potentially malign trading behaviour.

6 See the following paper for an example of trading algorithms colluding: Cartea, Álvaro and Chang, Patrick and Penalva, José, Algorithmic Collusion in Electronic Markets: The Impact of Tick Size (May 10, 2022). Available at SSRN: https://ssrn.com/abstract=4105954 *or* http://dx.doi.org/10.2139/ssrn.4105954

7 Wang, X.; Hoang, C.; Vorobeychik, Y.; Wellman, M.P. Spoofing the Limit Order Book: A Strategic Agent-Based Analysis. *Games* 2021, *12*, 46. https://doi.org/10.3390/g12020046 . Note: Although this paper is not specifically about machine learning based trading algorithms, the study shows that some trading strategies looking at order book information might be susceptible to be manipulated.

Risks	A more granular overview of risks of automated trading as we found in the literature and deemed relevant for this theme:
Trading algorithms programmed to manipulate	'In contrast to the analog, human protagonists of traditional market manipulation, new market manipulation generally uses the electronic communications, information systems, and algorithmic platforms of the new, high-tech financial marketplace to unfairly distort information and prices relating to financial instruments or transactions. At its core, these distortive actions and effects tamper with the humans and computerized information and communications systems of the marketplace. They corrupt how humans and machines communicate between and amongst each other in the financial markets. As such, this Article has termed this new approach to market manipulation, cybernetic market manipulation.' Pinging: 'With pinging, a larger number of small orders for a particular financial instrument are submitted and cancelled in fractions of a second by computerized platforms to induce others in the marketplace to react to their "pings" and disclose their trading intentions to the pinging party. Pinging allows the initiating party to discern valuable information at little to no risk since most of the pinging orders are cancelled prior to execution.' Spoofing: 'With spoofing, orders are placed by computerized platforms for a financial instrument at prices outside the current bona fide limits to spook other market participants to react in a manner favorable to the spoofing party. Spoofing allows the initiating party to distort the ordinary price discovery in the marketplace by placing orders with no intention of ever executing them and merely for the purpose of manipulating honest participants in the marketplace.' Electronic front running: 'Electronic front running is both similar and dissimilar from its traditional counterpart. Like its traditional counterpart, electronic front running seeks to manipulate the marketplace by executing trades ahead of a known future price change, thereby profiting once the price moving order is executed. Unlike its traditional counterpart that front ran traders via human brokers in small batches, electronic front running frequently leverages new, high-tech mechanisms that allow brokers to gain an unfair glimpse into order flows at one trading venue and to jump ahead of those flows to their advantage at another trading venue.' Mass misinformation: 'With mass misinformation schemes, parties can manipulate the marketplace through fake regulatory filings, fictitious news reports, erroneous data, and hacking. Because the new financial marketplace is so reliant on interconnected information and communications systems, a distortion to one source of information can have a large, volatile cascading effect on the greater marketplace in the short run, and a confidence-jarring effect on the greater marketplace in the long run.' *Lin, Tom C. W., The New Market Manipulation (July 3, 2017). Emory Law Journal, Vol. 66, p. 1253, 2017, Temple University Legal Studies Research Paper No. 2017-20, Available at SSRN:* https://ssrn.com/abstract=2996896
Machine learning and herding	'For example, AI models can identify signals and learn the impact of herding, adjusting their behaviour and learning to front run based on the earliest of signals. The scale of complexity and difficulty in explaining and reproducing the decision mechanism of AI algos and models makes it challenging to mitigate these risks.' *OECD (2021), Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges, and Implications for Policy Makers,* Artificial Intelligence, Machine Learning and Big Data in Finance - OECD

Risks	A more granular overview of risks of automated trading as we found in the literature and deemed relevant for this theme:
Explicit machine learning based trading algorithm learning to manipulate	'In this study, I constructed an artificial intelligence using a genetic algorithm that learns in an artificial market simulation, and investigated whether the artificial intelligence discovers market manipulation through learning with an artificial market simulation despite a builder of artificial intelligence has no intention of market manipulation. As a result, the artificial intelligence discovered market manipulation as an optimal investment strategy. This result suggests necessity of regulation, such as obligating builders of artificial intelligence to prevent artificial intelligence from performing market manipulation.' *Takanobu Mizuta, Does an artificial intelligence perform market manipulation with its own discretion? – A genetic algorithm learns in an artificial market simulation, (21 May 2020) 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*
Collusion	See the following paper about trading algorithms colluding: *Cartea, Álvaro and Chang, Patrick and Penalva, José, Algorithmic Collusion in Electronic Markets: The Impact of Tick Size (May 10, 2022). Available at SSRN: https://ssrn.com/abstract=4105954 or http://dx.doi.org/10.2139/ssrn.4105954* 'In digital marketplaces, increasingly sophisticated AI pricing agents (e.g. those based on DRL[8] methods) could discover, by self-learning, how to coordinate behaviors with their rivals, without being expressly instructed by their human developers or users, while also pursuing an optimal and rational strategy, like profit maximization. Under this novel scenario, independent AI traders employed by competing firms would be sufficiently sophisticated to self-learn the best policy actions and experiment with different strategies to optimize their (joint) cumulative performance. Therefore, "tacit" collusion would result from independent AI agents' autonomous decisions, without any prior human intent to achieve such a level of policy coordination.' *Azzutti, Alessio and Ringe, Wolf-Georg and Stiehl, H. Siegfried, Machine Learning, Market Manipulation and Collusion on Capital Markets: Why the "Black Box" Matters (January 6, 2022). European Banking Institute Working Paper Series 2021 - no. 84, University of Pennsylvania Journal of International Law, Vol. 43, No. 1, 2021* 1, 2021, Available at SRN: https://ssrn.com/abstract=3788872 or http://dx.doi.org/10.2139/ssrn.3788872
Herding, excess volatility, illiquidity, flash crash	'Similar to non-AI models and algos, the use of the same ML models by a large number of finance practitioners could potentially prompt of herding behaviour and one-way markets, which in turn may raise risks for liquidity and stability of the system, particularly in times of stress. Although AI algo trading can increase liquidity during normal times, it can also lead to convergence and by consequence to bouts of illiquidity during times of stress and to flash crashes. Market volatility could increase through large sales or purchases executed simultaneously, giving rise to new sources of vulnerabilities. Convergence of trading strategies creates the risk of self-reinforcing feedback loops that can, in turn, trigger sharp price moves. Such convergence also increases the risk of cyber-attacks, as it becomes easier for cybercriminals to influence agents acting in the same way. The abovementioned risks exist in all kinds of algorithmic trading, however, the use of AI amplifies associated risks given their ability to learn and dynamically adjust to evolving conditions in a fully autonomous way. For example, AI models can identify signals and learn the impact of herding, adjusting their behaviour and learning to front run based on the earliest of signals. The scale of complexity and difficulty in explaining and reproducing the decision mechanism of AI algos and models makes it challenging to mitigate these risks.' *OECD (2021), Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges, and Implications for Policy Makers,* Artificial Intelligence, Machine Learning and Big Data in Finance - OECD

8 Deep Reinforcement Learning

Risks	A more granular overview of risks of automated trading as we found in the literature and deemed relevant for this theme:
Systematic / macro risks	'The regulation and control of financial activity can be classified into two main categories, micro and macro. Micro control, to be executed by the micro AI, encompasses microprudential regulations and most internal risk management in financial institutions. It is inherently concerned with day-to-day activities of financial institutions, is hands-on and prescriptive, designed to prevent large losses or fraudulent behaviour, mandating and restricting how institutions should operate, what they can and cannot do, codified in the rulebook. While the rulebook was once in paper form, it is now increasingly expressed as digital logic, allowing programmatic access. Most, but not all, of the objectives a micro AI has to meet exist in the rulebook, and it generally has an ample number of repeated similar events to train on. All of this facilitates the application of AI to micro financial problems. Longer term objectives, such as the solvency of key institutions, financial stability and tail risk, risks that threaten the functioning of the financial system – systemic risk – are macro problems. Inside the regulatory space, that encompasses macro prudential regulations, and in the private sector, the management of solvency and liquidity risks for large financial market participants such as banks, insurance companies or mutual funds. The macro task is much harder. Macro risk is created by the strategic interactions of many players and involves aggregate phenomena such as bank runs or fire sales (Benoit et al., 2017).' *Danielsson, Jon and Macrae, Robert and Uthemann, Andreas, Artificial Intelligence and Systemic Risk (May 28, 2021). Journal of Banking and Finance* Journal of Banking and Finance, Forthcoming, Available at SSRN: https://ssrn.com/abstract=3410948 or http://dx.doi.org/10.2139/ssrn.3410948

Annex 3 **RTS 6 articles possibly affected by lack of explainability**

Take the following articles in RTS 6. This list is intended to be illustrative, not exhaustive:

- *RTS 6: article 2(1) Role of the compliance function: 'An investment firm shall ensure that its compliance staff* ***has at least a general understanding*** *of how the algorithmic trading systems and trading algorithms of the investment firm operate. The compliance staff shall be in continuous contact with persons within the firm who have detailed technical knowledge of the firm's algorithmic trading systems and algorithms'*

- *RTS 6: article 3(4): Staffing: 'An investment firm shall ensure that the staff responsible for the risk and compliance functions of algorithmic trading have: (a)* ***sufficient knowledge of algorithmic trading*** *and strategies; (b) sufficient skills to follow up on information provided by automatic alerts; (c)* ***sufficient authority to challenge staff*** *responsible for algorithmic trading where such trading gives rise to disorderly trading conditions or suspicions of market abuse.'*

- *RTS 6: Article 5(4): 'The methodologies referred to in paragraph 1 shall ensure that the algorithmic trading system, trading algorithm or algorithmic trading strategy: (a)* ***does not behave in an unintended manner****'*

- *RTS 6: Article 7 (1): 'An investment firm shall ensure that testing of compliance with the criteria laid down in Article 5(4)(a), (b) and (d) is undertaken in an environment that is separated from its production environment and that is used specifically for the testing and development of algorithmic trading systems and trading algorithms.'*

- *RTS 6: Article 13: Automated surveillance system to detect market manipulation*

Any questions or comments about this publication?

Send an email to: redactie@afm.nl

The Dutch Authority for the Financial Markets
PO Box 11723 | 1001 GS Amsterdam

Telephone
+31 20 797 2000

www.afm.nl

Data classification
AFM - Publiek

Follow us:

The AFM is committed to promoting fair and transparent financial markets.

As an independent market conduct authority, we contribute to a sustainable financial system and prosperity in the Netherlands.