THE CAPCO INSTITUTE
JOURNAL
OF FINANCIAL TRANSFORMATION

TRANSFORMATION

Artificial intelligence: Chances and challenges in quantitative asset management

FABIAN DORI | EGON RÜTSCHE
URS SCHUBIGER

DESIGN THINKING

#48 NOVEMBER 2018

THE CAPCO INSTITUTE

JOURNAL OF FINANCIAL TRANSFORMATION

RECIPIENT OF THE APEX AWARD FOR PUBLICATION EXCELLENCE

CONTENTS

DESIGN

TRANSFORMATION

SUPERVISION

DEAR READER,

Design thinking, a collaborative, human-focused approach to problem-solving, is no longer just for the creative industries. It has become an important management trend across many industries and has been embraced by many organizations. Its results are hard to ignore. Indeed, design-driven companies regularly outperform the S&P 500 by over 200 percent.[1]

To date, the financial services industry has not led in adopting this approach. However, leaders are recognizing that important challenges, such as engaging with millennial customers, can be best addressed by using design thinking, through the methodology's exploratory approach, human focus, and bias towards action. This edition of the Journal examines the value of design thinking in financial services.

Design thinking introduces a fundamental cultural shift that places people at the heart of problem-solving, which is critical in a technology-driven environment. If the customer's real problems are not fully understood, technological solutions may fail to deliver the desired impact. In this context, design thinking offers a faster and more effective approach to innovation and strategic transformation.

The case studies and success stores in this edition showcase the true value of design thinking in the real world, and how this approach is an essential competitive tool for firms looking to outperform their peers in an increasingly innovation-driven and customer-centric future. At Mastercard, design thinking has become a part of almost all organizational initiatives, from product development, research and employee engagement to solving challenges with customers and partners. Meanwhile, at DBS Bank in Singapore, a data-informed design model has been firmly embedded into the bank's culture, enabling them to successfully move from being ranked last among peers for customer service in 2009, to being named the Best Bank in the World by Global Finance in 2018.

I hope that you enjoy the quality of the expertise and points of view on offer in this edition, and I wish you every success for the remainder of the year.

Lance Levy, **Capco CEO**

1 http://fortune.com/2017/08/31/the-design-value-index-shows-what-design-thinking-is-worth/

ARTIFICIAL INTELLIGENCE: CHANCES AND CHALLENGES IN QUANTITATIVE ASSET MANAGEMENT

FABIAN DORI | Quantitative Strategist, AQ Investment Ltd.
EGON RÜTSCHE | Quantitative Strategist, AQ Investment Ltd.
URS SCHUBIGER | Quantitative Strategist, AQ Investment Ltd.

ABSTRACT

Artificial intelligence has recently experienced a remarkable increase in attention, following staggering achievements in applications such as image, text and speech recognition, self-driving cars, or chess and Go tournaments. It is, therefore, not surprising that the financial services industry is working hard to improve investment decisions by incorporating self-learning algorithms into the investment process. However, do all sectors of the asset management industry exhibit characteristics that can benefit from applying artificial intelligence tools to uncover new patterns? Are there limits beyond which additional computing power and greater data availability have only marginal benefits? This article provides some initial answers to these questions. It demonstrates that the adaptivity and self-learning capability of machine learning tools could add value along the entire value chain of an asset manager. However, the inherently flexible nature of machine learning methods is also the biggest challenge. These methods must be applied thoughtfully and in the right context. The article provides a general overview of machine learning, elaborates on specific applications in quantitative asset management, and highlights limitations, challenges and possible remedies.

Artificial intelligence has in recent years received a lot of attention, following staggering achievements in various applications. It is, therefore, not surprising that the financial industry is also trying to improve investment decisions by incorporating self-learning algorithms into the investment process. Quantitative tools and algorithms have been used within the hedge fund industry to define systematic trading strategies for some time now, which is why quantitative hedge funds may provide a fertile ground for the application of new machine learning techniques. This article shows that the adaptivity and self-learning capability of machine learning tools may add value along the entire value chain of an asset manager. However, the inherently flexible nature of machine learning methods is also their biggest challenge. It requires that the methods are put in the right context and thoughtfully applied. This article begins with a general overview of machine learning, then elaborates on specific applications in quantitative asset management, also highlighting limitations, challenges, and possible remedies before concluding in summarizing remarks.

1. FROM MACHINE LEARNING IN GENERAL…

Machine learning refers to extracting knowledge from data by identifying correlated relationships without getting prior information about what causal dependencies to look for. It combines elements from both statistics and computer science and has been in existence for many years. However, it has been mostly due to significant advancements in computing power and data availability that the application of artificial intelligence algorithms has become applicable for everyday life in recent years.

Most machine learning methods have been developed outside of finance and built on well-known statistical models, such as linear regression or clustering techniques. Still, the machine learning framework allows for much more flexibility. It can be applied to different kinds of problems, such as classification or regression analysis. Classification algorithms group observations into a finite number of categories, whereas regression analysis estimates outcomes of problems that have an infinite number of solutions. While machine learning is a very broad field, it can be classified into three main areas (Table 1).

The most successful field currently is **supervised learning**, where algorithms learn based on provided training data that reveal known relationships. The simplest form of a supervised learning algorithm is **linear regression**, which makes a prediction using a linear function of the input features. There is a general trade-off between optimizing the fit of a model on the in-sample training and the true out-of-sample prediction period. Given that all models tend to fit the training data better when more input variables are used, it may be reasonable to penalize additional model complexity in order to maintain sufficient generalization power for the prediction task. Methods such as "ridge" or "lasso regression" help in automatically detecting the most relevant input variables by regularizing model complexity to avoid overfitting. All three methods are by their nature linear, but may account for nonlinear relationships based on an appropriate manipulation of the input variables. A simple machine learning method that is not constrained to linear relations is the "k-Nearest Neighbors algorithm". This model looks for the k historical data points that come closest to the current situation and predicts future values based on these historical "neighbors". There are more complex nonlinear supervised learning algorithms, such as **decision trees** or **random forests**, however, they are not able to extrapolate and to make forecasts outside of the range of the training data.

Contrary to the methods described above, **unsupervised learning** algorithms only receive input data to learn from, but no information about the output data or relationships. These algorithms, therefore, detect patterns in the data by identifying clusters of observations that depend on similar characteristics. Machine learning can, for example, be used to identify the main topics in the news flow for a given stock. At the core of unsupervised learning algorithms is the idea of reducing dimensionality by clustering the data or by transforming it into simpler factor representations. **Clustering** methods partition the input data into subsets that exhibit common characteristics, such that the data points within a cluster share some notion of similarity that decisively separates them from the data points in other clusters. **Factor analysis**, on the other hand, relies on transforming the original data into the most relevant drivers or the most appropriate representation.

Combining methods of supervised and unsupervised learning results in the so-called **reinforcement learning**, where the algorithm first detects patterns on its own, and then receives feedback from an exogenous source to validate or further guide the learning process. A reward feedback is required for the algorithm to learn a certain behavior.

The artificial intelligence literature also frequently refers to **deep learning** or **neural network** algorithms. This kind of method mimics, in a certain sense, the functions of the human brain by feeding information through different layers and nodes. The simplest form is called "multilayer perceptron" and can be seen as a generalization of linear models that perform multiple regression steps. There exist more advanced networks to deal with the challenges of simple networks and to allow for greater complexity. Given that even a sketchy synopsis of this area would be beyond the scope of this article, we refer the interested reader to the corresponding abundant literature.

Table 1: Artificial intelligence and exemplary methods

<table>
<tr><th>SUPERVISED LEARNING</th><th>UNSUPERVISED LEARNING</th></tr>
<tr><td>Linear regression, Ridge, Lasso,
k-Nearest Neighbors, decision trees</td><td>Clustering (k-means),
Factor analysis (PCA, manifold learning)</td></tr>
<tr><td colspan="2">REINFORCEMENT LEARNING</td></tr>
<tr><td colspan="2">DEEP LEARNING</td></tr>
<tr><td colspan="2">Multilayer, feed-forward neural networks</td></tr>
</table>

Table 2: Artificial intelligence applications in quantitative finance

INVESTMENT PROCESS	POSSIBLE MACHINE LEARNING APPLICATIONS	EXEMPLARY METHODS
INVESTMENT UNIVERSE	• Identifying uncorrelated assets • Data mapping allowing for interpretation • Proxy assets with more liquid instruments	• Dendrogram • PCA, manifold learning • k-means
SIGNAL CREATION	• Create nonlinear forecast • Forecast based on predefined factors • Forecast without prior factor knowledge • Forecast based on new data sources	• Linear, lasso, ridge regression • k-nearest neighbors • Bag-of-words, term frequency-inverse document frequency
PORTFOLIO CONSTRUCTION	• Improve estimate of input parameters • Dynamically maximize target value • Scenario and stress testing	• PCA • Lasso, ridge regression • k-means
TRADING	• Pricing based on sparse data • Non-linear relations in transaction costs	• k-means • Lasso, ridge regression

2. …TO SPECIFIC APPLICATIONS IN QUANTITATIVE ASSET MANAGEMENT…

Each sector within the financial services industry uses artificial intelligence methods differently and according to their specific needs. In what areas of asset management can artificial intelligence provide added value? And what problems of investment managers may not be solved by such tools?

When we think about modeling the investment process – and hence about using machine learning algorithms to improve decision making – we like to subdivide the value chain into different steps. This allows for a systematic application of models that are appropriate for the specific task. We will provide details for the various steps in the next section, but start with an overview in order to facilitate the synopsis (Table 2). The first step consists of the "definition the investment universe". Next, the "alpha engine" or "signal engine" preprocesses the data in a proper way, calculates the signals for the various markets under scrutiny based on the models used, and maps these signals into the portfolio context. Subsequently, the "portfolio construction engine" or "risk management engine" builds the theoretical model portfolio based on a suitable algorithm, taking regulatory and investor specific limits into account. Finally, the "trading engine" translates changes in the model portfolio positions into effective trades.

What then, do these more general descriptions adhere to in more detail? Let us start with the investment universe, where machine learning tools may add value by identifying uncorrelated assets that provide true diversification benefits, or by mapping data into new representations that allow for other interpretations, such as the detection of style drifts in hedge fund strategies or factor exposures such as momentum or value. An appropriate tool for the first task would, for example, be a dendrogram analysis (see case study 1: Clustering the investment universe with dendrograms). The second goal may be achieved by relying on a principal component or manifold learning analysis. In a similar manner, artificial intelligence methods can be used to proxy the valuation or even actual investment of assets that have only sparse historical market data or are not eligible due to liquidity issues with more liquid instruments that appropriately mimic the characteristics of the actually desired assets. A useful variant to achieve that task would be the k-Nearest Neighbor model.

Case study 1: Clustering the investment universe with dendrograms

Dendrograms belong to methods of hierarchical clustering. The algorithm iteratively clusters first individual data points and then sub-clusters into a hierarchical order, depending on their respective correlation structure.

We use a dendrogram to structure a set of individual commodity markets into more meaningful clusters. Ideally, it comes up with the well-known sectors such as energy, precious and industrial methods for example.

At the bottom of the visual representation in Figure 1 are the single data points that are joined in first clusters. For example, the model groups copper and aluminium to a mutual cluster of industrial metals, or gold and platinum to precious metals. These two clusters are then joined to form a cluster of metals in more general. Similarly, heating oil and crude oil are merged before being clustered together with gas oil as the energy complex. The energy and metals cluster are then put together to form a cluster of commodities that are more dependent on business cycle swings. The soft commodities soybeans, soybean meal, corn, and wheat are structured in a separate node that only consists of agricultural products. Interestingly, the natural gas commodity forms an individual cluster, most likely because of seasonality factors that separate it from the other energy commodities.

Figure 1: Dendrogram for investment universe

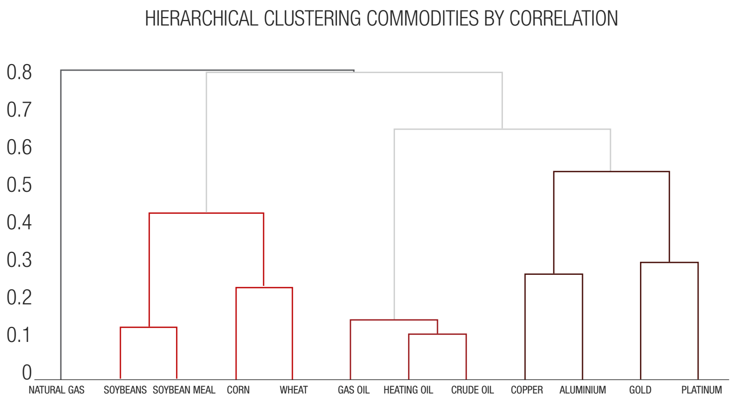

Figure 2: KNN for predictive power of moving averages

The aim of the alpha or signal engine is to produce forecasts on the direction and magnitude of future price movements of the relevant assets or their respective riskiness, and how to translate that information into a meaningful signal for the portfolio construction engine. Potential applications of machine learning methods in this field can be classified into three main blocks. First, an artificial intelligence algorithm may be helpful to create a nonlinear forecast based on a single time-series (see case study 2: Analyzing the behavior of the VIX index with KNN). Second, machine learning methods may derive the forecast value out of a predefined pool of relevant factors. More involved is the third application, consisting of letting the model select relevant input signals on its own or access new data sources in order to extract additional information.

Case study 2: Analyzing the behavior of the VIX index with KNN

The VIX index measures the market expectations for the volatility of the S&P 500 index over the coming month, based on index option prices. Given that volatility can neither become negative nor grow boundlessly, theory suggests a mean reverting behavior. Additionally, the distribution of changes in volatility levels is commonly skewed, mimicking the fact that spikes in volatility oftentimes occur very fast, while the slowdown in volatility normally takes more time and is a bumpier road.

Given that backdrop, we analyze the predictive power of two moving averages on past index movements for the future direction of the VIX Index based on a k-nearest neighbor classification algorithm. Figure 2 plots the values of the short moving average on the x-axis and the values of the long moving average on the y-axis. Conditional on the value of these two moving averages, the blue points represent moments in time where the future VIX index movement was negative, and the green points indicate future positive directional changes. While it is difficult to extract a meaningful interpretation from this scatter plot, a k-nearest neighbor analysis reveals further information. Based on this estimator, Figure 2 shows the decision boundaries for the two states of future directional movements in separate colors. The area colored in blue represents states where the two moving averages indicate falling VIX levels, whereas the red area stands for scenarios in which the two moving averages predict a rising VIX. Clearly, positive values for the moving averages are related to negative future VIX price movements, confirming the mean reverting behavior after an increase in volatility levels. The picture for negative moving average values is more ambiguous, overall upholding the thesis of mean reversion, but also showing some signs of momentum. That makes intuitive sense, as volatility tends to trend lower after a sudden outburst.

> "Artificial intelligence aims to extract relevant knowledge from possibly unstructured data on a self-learning basis."

Based on these forecasts, the portfolio construction or risk management engine calculates the target positions, taking regulatory and internal restrictions into account. In this step, artificial intelligence methods may be helpful to improve the estimate for input variables. This could, for example, be achieved by reducing the dimensionality of the dataset based on clustering algorithms such as a principal component analysis. Instead of optimizing the portfolio with respect to a predefined objection function and specific constraints, machine learning tools may also be asked to tweak the portfolio in a more general way. For instance, by dynamically weighting the portfolio components such that risk-adjusted returns in the sense of Sharpe ratios or the ratio of average returns to maximum drawdown are maximized. Finally, enhanced scenario analysis tools may improve model validation and stress testing applications.

The trading engine finally translates the target positions into effective market orders. This step is especially relevant for large asset managers, as an estimated two-thirds of gains on trades are lost due to market impact costs when trading into and out of large position blocks.

3. …TO CHALLENGES AND LIMITATIONS

In the previous section, we have highlighted various steps along the value chain of a well-structured investment process that in our view are suitable to be further enhanced by machine learning applications. However, there are a number of challenges and limitations that are not necessarily new to quantitative investment managers, but may be aggravated by the flexibility of new techniques.

Artificial intelligence aims to extract relevant knowledge from possibly unstructured data on a self-learning basis. It works especially well for tasks with precisely defined rules and stable probability distributions, such as mastering demanding games like chess or Go. However, the stochastic nature of financial markets with its lack of stable rules and probability distributions may challenge the validity of relationships that are learned from the past. Accordingly, models should always be applied to clearly defined problems and validated against sound theoretical assumptions.

Similarly, self-driving cars can be driven along the same roads as many times as is necessary to teach them all the relevant aspects and AlphaGo can play with itself until it perfectly masters the rules of the game. However, despite the seemingly abundant access to data, there is only one historical price trajectory for each financial market to train a model on. This limited data availability restricts the complexity of the artificial intelligence model that can be applied and, therefore, the flexibility of its output when forecasting future price movements. This problem is further exacerbated by the fact that the vast majority of data for financial markets has only been collected recently. As a consequence, researchers should focus on parsimonious model structures and not be misled by the mightiness of artificial intelligence models to adaptively learn the past.

Table 3: Challenges and possible remedies

CHALLENGES FOR ML IN ASSET MANAGEMENT	POSSIBLE REMEDIES
LACK OF STABLE RULES AND PROBABILITY DISTRIBUTIONS	• Apply models to clearly defined problems • Validate results against sound theoretical assumptions
LIMITED DATA AVAILABILITY	• Focus on parsimonious model structures • Account for mightiness of models to learn the past
LACK OF CAUSAL REASONING AND IMAGINATION	• Keep simplified transformation in mind • Diligent analysis before investing resources and modelling power
FUTURE REGULATION AND SUSCEPTIBILITY TO MANIPULATION	• Check future accessibility of data sources • Validate contributors to data sources

Next, machine learning models excel at identifying relationships in the data that may be unrecognizable to the human eye. Still, they lack causal reasoning and imagination that would be necessary to anticipate events that have not happened in the same way many times before. Would a trading model based on artificial intelligence have been able to predict the currency peg break between the Euro and the Swiss Franc by the Swiss National Bank in early 2015? Most likely not. In a similar manner, machine learning algorithms may just find theories that are already well-known and proven. While this confirmation may add comfort, it may also just be a waste of time and money. So, despite increasing computer power and data availability, it is necessary to keep in mind that quantitative models remain a simplified transformation of the world and will only have forecasting ability that is limited to specific tasks. Additionally, the complexity of calibrating artificial models requires a diligent analysis on where to allocate resources and model power most effectively.

Other more general potential limitations include future regulation and the susceptibility to manipulation.

4. CONCLUSION

This article provides a framework to assess the opportunities and challenges of applying artificial intelligence methods within a structured investment process. It highlights that the adaptivity and self-learning capability of machine learning tools may add value along the entire value chain of an asset manager. First, by more effectively using currently available data based on algorithms that learn to reveal new nonlinear relations or transform it into representations with more interpretable meanings. Second, by embracing new data sources that provide additional information. However, the inherently flexible nature of machine learning methods is also their biggest challenge. It requires that the methods are put in the right context and thoughtfully applied to solve questions that produce meaningful outcomes. It would be illusionary to believe that artificial intelligence will develop a profitable investment rationale on its own. Accordingly, we are convinced that machine learning will most likely turn out not to be the much-searched Holy Grail, but that it will help quantitative investment managers in further improving their allocation processes. Nevertheless, the application of artificial intelligence in asset management is still in its early days. This paper, consequently, provides evidence on first experiences, but no final results. We are looking forward to an exciting future.

THE FUTURE. NOW.

ABOUT CAPCO

Capco is a global technology and management consultancy dedicated to the financial services industry. Our professionals combine innovative thinking with unrivalled industry knowledge to offer our clients consulting expertise, complex technology and package integration, transformation delivery, and managed services, to move their organizations forward. Through our collaborative and efficient approach, we help our clients successfully innovate, increase revenue, manage risk and regulatory change, reduce costs, and enhance controls. We specialize primarily in banking, capital markets, wealth and investment management, and finance, risk & compliance. We also have an energy consulting practice. We serve our clients from offices in leading financial centers across the Americas, Europe, and Asia Pacific.

To learn more, visit our web site at **www.capco.com, or follow us on Twitter, Facebook, YouTube and LinkedIn.**

WORLDWIDE OFFICES

APAC
Bangalore
Bangkok
Hong Kong
Kuala Lumpur
Pune
Singapore

EUROPE
Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
London
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA
Charlotte
Chicago
Dallas
Houston
New York
Orlando
Toronto
Tysons Corner
Washington, DC

SOUTH AMERICA
São Paulo

WWW.CAPCO.COM