

Artificial intelligence in investing

Great aspirations – reverie or reality?

About the author



Daniel Seiler, PhD
Head of Multi Asset

Daniel Seiler heads the Multi Asset Boutique at Vontobel Asset Management. In this role, he oversees systematic and fundamental investment solutions, which are offered by the Vescore and Vontobel brands respectively. He is responsible for the investment processes, driving innovation in product development and managing the investment teams.

Previously, from 2017 to 2018, he was Head of Vescore, the quantitative investing franchise of Vontobel Asset Management. Daniel Seiler became Chief Investment Officer at Vescore Solutions AG in 2009 and was instrumental in operating the company's portfolio and investment management division.

Prior to this, he ran the hedge fund platform of Swiss Capital Investment AG, where he oversaw the investment activities of various funds of hedge funds. Daniel Seiler started his career as an analyst in the field of sustainable investments.

Daniel Seiler's educational background includes a PhD (Dr. oec. HSG) and a Master's degree in Financial and Capital Markets Theory from the University of St. Gallen (lic. oec. HSG). In addition, he holds a Master's degree in Environmental Sciences from ETH Zurich (dipl. natw. ETH).

Edition: March 2020



Vontobel uses only recycled paper for printing. It takes about 1.5 times less energy and 2.5 times less water to produce recycled paper than it does to produce paper from fresh fiber. Recycled paper also cuts greenhouse gas emissions by more than 20%. We offset the remaining emissions with various CO₂ projects around the world.

Further information:
vontobel.com/sustainability

Content

4

Grasping the puzzle of intelligence

The emergence of artificial intelligence

Artificial intelligence stripped down to its essentials

6

Machine learning

Deep learning

7

Neural networks under the magnifying glass

8

Data shortage might throw a spanner in the works of AI

9

Unlocking value with artificial intelligence in financial markets

Artificial intelligence joins the financial industry's chase for predictability

10

How can AI add value?

11

How AI became the game master of chess

18

When can AI add value?

20 – 21

AI's stumbling blocks and stepping stones in finance

New techniques in the artificial intelligence toolkit harbor potential pitfalls

The black box problem

22 – 23

Our ability to tell stories differentiates us from machines

The white box solution

The future of artificial intelligence in investing is bright

Executive Summary

Daniel Seiler, PhD

Head of Multi Asset

Investing is the ultimate numbers game and smart number crunchers tend to be good at it. So, artificial intelligence (AI) as a high-capacity data processor stands a good chance at revolutionizing the investment industry. Over the past few decades, AI has risen to new prominence and expanded its capabilities on the back of the exponential rise in computing capacity and data volumes. Thanks to its potent enablers machine and deep learning, AI has encroached on many areas of human activity where it at least matches or even surpasses human intelligence.

Achieving superior performance in financial markets poses particular challenges. AI's intelligent algorithms have the ability to place more and better bets on markets in a cost efficient way. This allows them to obtain a certain degree of predictive power on the fundamental factors that move asset prices. Competitive forecasting power is any market participant's greatest asset when investing with success. Certain market structures must be present for AI to operate to the full extent of its capabilities. Only if markets are less than perfectly efficient, can AI fully harness its powerful potential to act on information that is not yet reflected in asset prices.

However, AI engineers must address their algorithms main weakness lest they get stopped in their tracks: lack of transparency. So far, the inner workings of most sophisticated AI algorithms' remain a mystery not only to their end-users but also to their developers. This is because they rely on self-learning algorithms that generate highly accurate outputs despite close to no human intervention. Black box algorithms whose reasoning remains obscure are unlikely to establish credibility and trust with investors. In addition, they will have to wrestle with regulators who are alarmed by the opacity of the investment decisions they make. The remedy to this deficiency lies with crafting AI powered communication interfaces that convert the mathematical functions of an intelligent investment algorithm into stories with intelligible information for investors. Once it has shed its unfathomable nature, AI will be able to shrug off looming threats from regulators and have the potential to fuel credible and impactful investment strategies. Transparent AI investment engines create true value for investors. They turn unstructured data pools into investment insights and improve investment management forecasting methods. This results in an irresistible selling proposition that could attract long-lasting investor commitments.

Grasping the puzzle of intelligence

I place my king on H7 in an attempt to escape “check”. Swiftly, Tov Kronsteen takes his queen, lifts her up and stops in midair looking at me with derision. My stomach contracts. He puts her on E4, leans back and, bored by my inferiority, with a cigarette dangling from the corner of his mouth, he says “check”. My eyes scan the board. I feel blood rushing to my face. There are no more options. With a shaking hand I overturn my king indicating defeat. “Congratulations, Sir. A brilliant coup,” with these words Kronsteen shakes my hand with an air of deep satisfaction before he strides from the room, power emanating from his demeanor.

The above recounts the opening scene from the famous James Bond movie “From Russia with Love” which shows the villain Tov Kronsteen beating the Canadian citizen MacAdams at a chess duel. Kronsteen acts as the Director of Planning to James Bond’s main antagonist, Spectre, and boasts of an immense intellect while being an expert at scheming and plotting. At the end of the scene, he emerges as the victor and is celebrated by the audience.

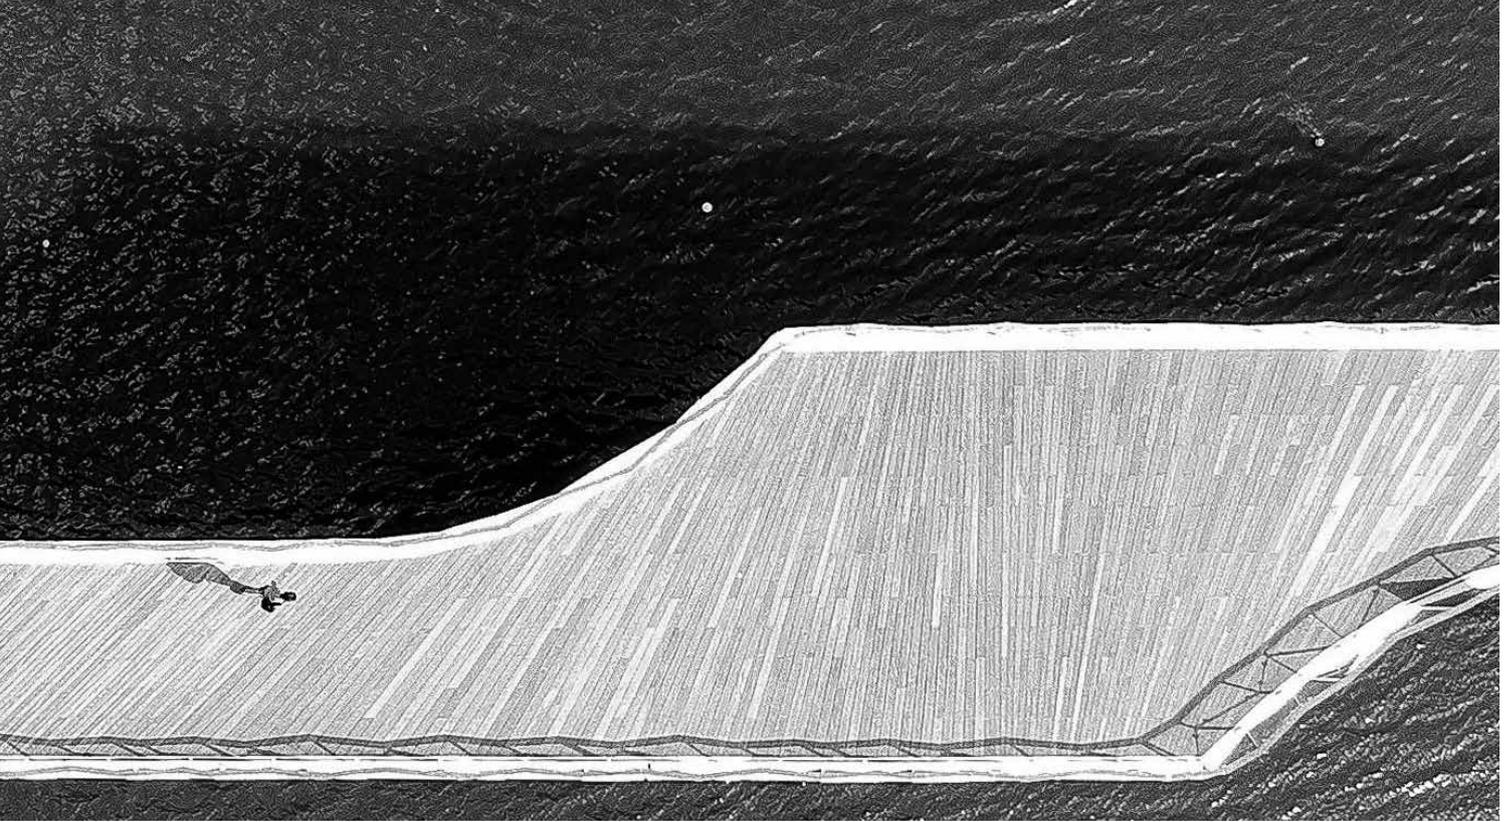
The movie, released in 1963, illustrates a conception of human intelligence which revolves around an individual’s intellectual and analytical prowess, calculating abilities and mental capacity which, together, form a genius-level intellect. Mastery of chess, one of the most demanding strategic board games of our times, is the apogee of human intelligence. Today, this idea seems outdated as intelligent computers have overtaken humans thanks to their superiority in number crunching and data processing (for more details see pages 11–16).

If chess is a numbers game, then so is investing. Mastering the financial markets game goes hand in hand with making money from it, an ability commonly associated with

smart people. So, applying artificial intelligence to financial markets seems alluring and begs the question if intelligent machines could outshine humans in this area too.

The emergence of artificial intelligence

Machines infringing on human abilities cast doubts on common conceptions of human intelligence which has been a puzzle to the brightest of philosophers and scientists throughout time. The last century alone produced a large variety of definitions trying to capture the essence of human intelligence. To name but a few, in 1904, the English psychologist Charles Spearman suggested that there was a general intelligence factor which makes people score high on IQ tests. Either you have it or you don’t. In 1983, Harvard psychologist Howard Gardner considered Spearman’s theory too simple a concept and proposed his Theory of Multiple Intelligences which claims the existence of eight distinct types of intelligence. A person could possess few or many of them. Gardner’s main critic, Robert Sternberg took a more cognitive approach and professed that intelligence can be broken down into analytic, creative and practical subsets.



More recently, neuroscientists have tried to understand why some brains are smarter than others by dissecting the neurobiological design of people's brains of different IQs.

If the debate of the last century seems heated, it is no comparison to the bewilderment and exacerbation that has befallen today's discussion which is far away from settling on a universally recognized definition. One of the culprits for this development are intelligent machines that challenge traditional notions of intelligence by calling into question what really makes the human mind different from other forms of intelligence.

The field of artificial intelligence first emerged in the 1950s when humankind progressed in its understanding of how the human mind works and applied scientific findings to making machines that emulate human intelligence.¹ However, only over the past few decades has this branch of information technology unfolded to its full glory as technological advances in computer processing power have made its tools attractive to corporate giants of various industries. One of the field's most powerful enablers, machine learning, has opened the door to smart algorithms making inroads in various areas of human life and activity. Self-driving cars are an obvious example. Social media platforms use facial recognition software when you upload pictures in order to suggest which people in the picture you might know. Product recommendations given by online stores are not chosen at random but highly consumer-specific with the goal to increase sales. The brain behind this? Artificial neural networks that learn what customers like. Other AI applications are usually taken for granted, like email spam filters that continuously learn which emails to block and which ones to let pass. Less profane examples are robo-readers grading your university exams, robo-personal assistants booking your flights

and robo-baristas predicting what you want to order in coffee shops before you even know you want it.

For the future, AI is full of promises. Many hope to eliminate burdensome, repetitive jobs and increase operational efficiency for their firms by smart automation processes. It is estimated that up to half of the activity in the economy has automation potential by means of robotics, AI and machine learning.² Many go even a step further by looking to AI as empowering a second industrial revolution. They expect the engines of the first industrial revolution, which liberated humankind from the yoke of manual labor, to be replaced by smart machines that change how modern human labor works by enabling existing engines to perform new tasks.

Artificial intelligence stripped down to its essentials

AI raises aspirations and encourages dreams. Unsurprisingly, it has become a global buzzword for many, lightly thrown about in various contexts. This has distorted its meaning. A definition will set the record straight.

“AI raises aspirations and encourages dreams.”

Artificial intelligence is a field of computer science that strives to emulate human intelligence in machines. Essentially, machines are trained to perform tasks that ordinarily require human intelligence, such as speech recognition, story-telling, visual perception and decision-making. It is a multi-faceted field as it has a number of layers that increase in sophistication and complexity with the depth of learning achieved by the machine.

¹ AI is a term coined in 1955 by the American computer scientist John McCarthy who dedicated his work to developing intelligent machines.
² Boughin, J., Manyika, J. and Woetzel, J. (2017). A future that works: automation, employment and productivity. [ebook] McKinsey Global Institute. Available at: https://www.mckinsey.com/mwg-internal/de5fs23hu73ds/progress?id=6TfTsrwVViPeNn1ecR_ebKTiAPjXO-wNzVaS-WZPq1z0,&dl [Accessed 10 Feb. 2018].

Machine learning

Machine learning is an enabler of AI. It trains a computer how to perform a task or generate a predefined desired output without being specifically programmed for it. Instead, it combs large amounts of data in order to extract relevant aspects of the data, cluster and classify it. Most machine learning algorithms tend to be “supervised” which means that humans have to provide defined input data, known as features, as well as the desired output. In contrast, with “unsupervised” learning, the user only provides the features and the algorithms train themselves to fulfill objectives.

An artificial neural network is one of many computational algorithms used for the purpose of machine learning. Built with the intention to simulate biological neural systems, artificial neural networks come closest to the concept of an artificially engineered intelligence among all AI algorithms. In biological brains, electrochemical impulses transmit through a network of neurons that consist of a cell body, dendrites and axons. Dendrites are receptive extensions of a nerve cell that transport signals into the cell body which accumulates signals until a certain threshold is reached. Then, axons carry the electrochemical impulse away from the cell body to other neurons, muscles and glands.

Artificial neural networks emulate this behavior. While they do not match biological neural systems in complexity, they also consist of a web of artificial neurons, so called nodes, which are bundled in different layers that are interconnected. The nodes contain mathematical activation functions, configured to convert an input signal into an output signal. The learning done by these nodes can be understood as parametrization processes that take place in each activation function in every node.³ The machine tunes and tweaks a large amount of parameters until the input features received are converted into the desired output. Once learning is complete, the machine is able to apply its knowledge to unknown data. The artificial brain literally learned how to perform a task.

An example for such an undertaking is training a machine for simple categorization tasks. An exercise as profane as distinguishing an apple from an orange first requires the extraction of features such as shape, surface texture

and weight of these types of fruits. These features are subsequently fed as inputs to a simple neural network that consists of one node, three inputs and one single output (see illustration on page 7).

In addition, the desired outcome, an apple or orange, is defined. Learning, or more precisely the parameterization of the function inside the node, continues until a statistically significant amount (apples and oranges) has been identified and labeled correctly.

Deep learning

Deep learning takes machine learning one step further by performing more complex processing tasks by mimicking biological brains more closely. A human brain consists of approximately 100 billion neurons with each having around ten thousand connections.⁴ In order to approach this complexity, deep artificial networks consist of a high number of nodes and layers. A basic artificial neural network has one or two layers, whereas a deep network might have several dozen (see illustration on page 7). The key additional advantage between machine and deep learning is the ability to perform feature extraction without any human interference thanks to increased network depth. This means that the algorithm is able to derive meaningful features from the input data. Moreover, a deep network can digest enormous amounts of input data and process them through multiple layers which enables it to learn increasingly more complex data representation with each layer.

These additional capabilities come at a price, however. In order to perform well, deep neural networks depend on input data that is highly granular in nature. In addition, the sample data should be fairly consistent and not too “noisy” in order to enable a straight-forward training process and accurate outputs. The noisier the data, the more data is required. To a certain extent, the requirement for large data volumes can be met by stacking several pre-trained neural layers from multiple sources which allows for transferring knowledge from proven algorithms to new systems. This would not be possible with machine learning that is based on simple neural networks of insufficient depth. The results are applications driving remarkable advances in image, text and speech recognition.

→ continued on page 8

³ Parametrization is a mathematical process of defining parameters of a function.

⁴ Goodfellow, I., Bengio, Y. and Courville, A. (n.d.). Deep learning.

⁵ Digital images are pictures that are stored on computers in the form of a matrix of numbers that the computer can understand. This is why digital images lend themselves easily to training AI algorithms.

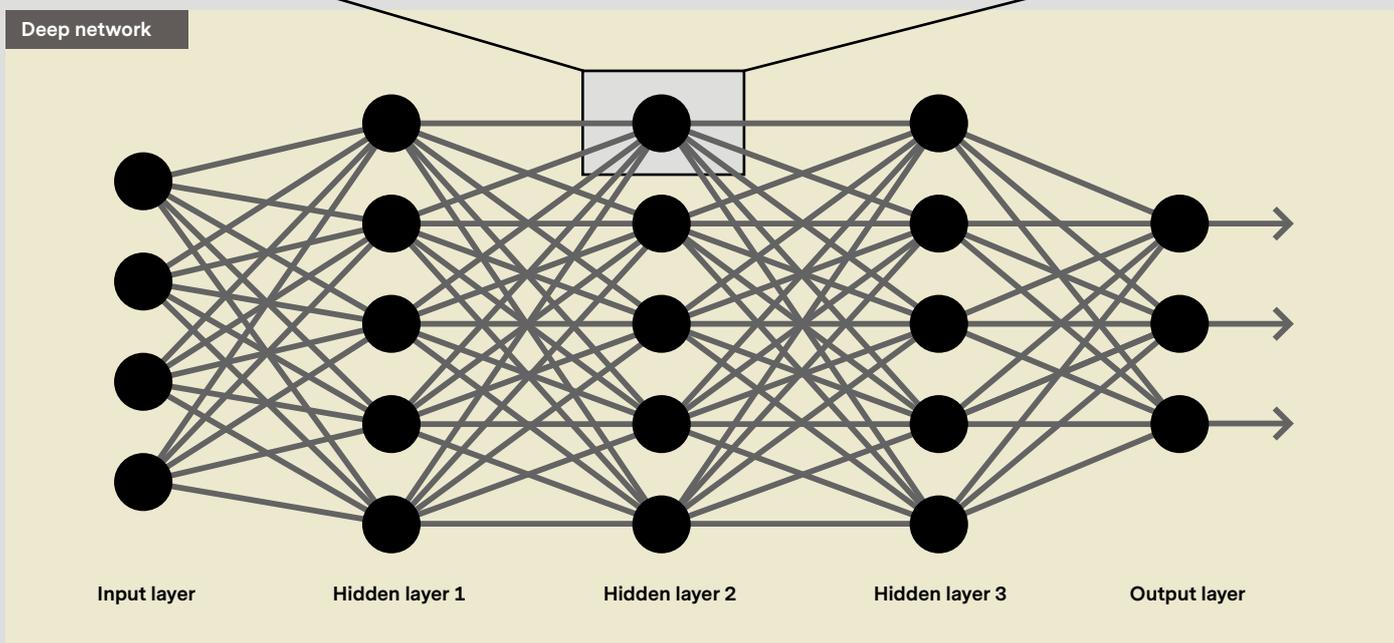
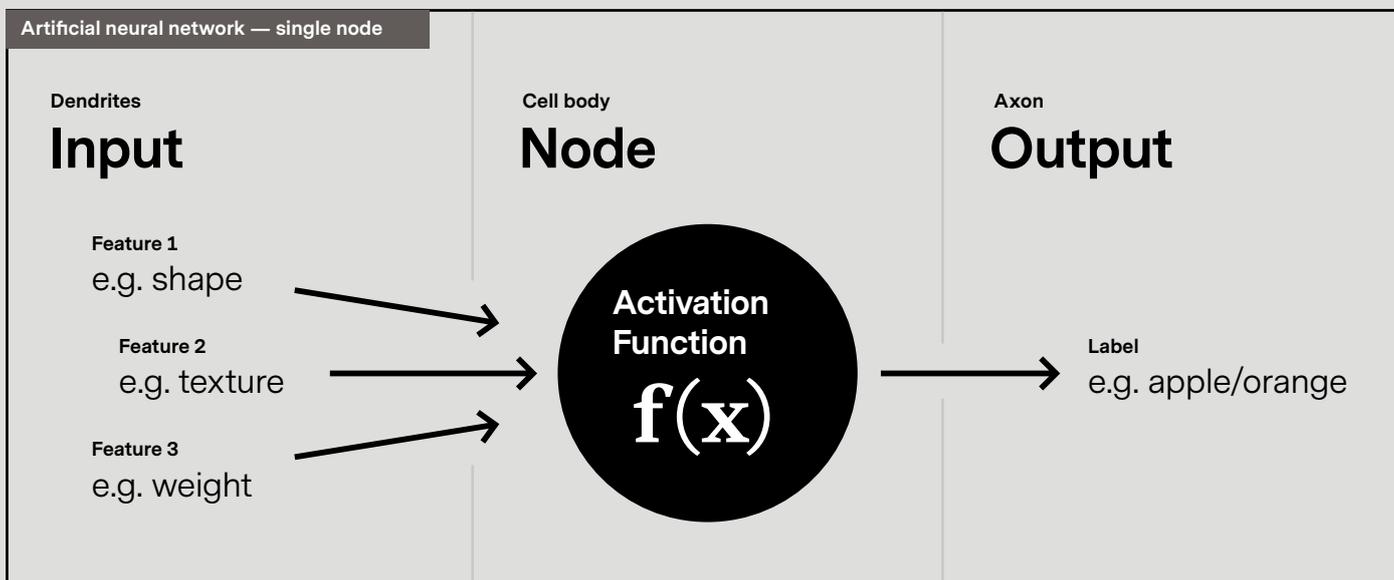
Neural networks under the magnifying glass

Neural networks are AI algorithms that mimic biological brains with varying degrees of complexity. The first illustration reveals the apparatus of a single node. This type of algorithm is used in basic machine learning and is able to perform simple categorization tasks such as distinguishing apples from oranges.

of networks have the ability to execute more complex tasks such as detecting and identifying objects against dynamic backgrounds. They are applied in deep learning approaches and are put to use in various areas of human economic activity. For example they are used in robotic farming which requires accurate fruit and vegetable detection in changeable environments such as farmland.

The drawing below showcases a deep neural networks that consists of a high number of nodes bundled in multiple stacked hidden layers. Due to their depth, these kinds

To learn more, please refer to the sections on machine and deep learning on page 6.



Let's continue the fruit categorization example from above. The application of a deep neural network to the task automates the cumbersome feature extraction process. Once trained with a large amount of labelled input pictures,⁵ the deep neural network by itself comes up with the required input features as well as the correctly labeled output data, i.e. "apple" or "orange". The ability to perform automated feature extraction has contributed considerably to the current "Industrial Revolution 2.0" phenomenon since highly complex processing tasks can be automated in various sectors of the economy. For example, robotic harvesting relies on the reliable, fast and precise classification of fruits in a complex environment such as farmland and fields. Fruits not only have to be located, they also have to be distinguished from a multitude of other objects such as leaves, stems and branches on a live video stream that is fed to the harvesting machine. This is further complicated by differing fruit

features such as color and shape as well as varying degrees of ambient lighting which might result in fruits blending in with the visual appearance of the background. The process of feature extraction from a dynamic pixel stream that is precise enough in order to distinguish fruits from a multitude of backgrounds can only be performed by a machine and not the human eye.

“Predictive power is what all asset managers seek, be they conventional active ones or AI driven quants.”

Data shortage might throw a spanner in the works of AI

As sophisticated and impressive as this sounds, the requirement of immense data volumes for deep artificial networks to work accurately might prevent them from being applied to domains that do not boast of huge consistent data repertoires, such as financial markets. This might seem counterintuitive at first. Because if investing isn't the ultimate numbers game, then what is? While it is true that a lot of number crunching goes into explaining and predicting market movements, the actual amount of data at hand that is suitable for financial modeling and analysis is often limited. Structural changes in underlying economic forces and political regimes as well as corporate actions such as mergers and acquisition may result in erratic data behavior and inconsistent long-term time series. Think about it this way, the economic conditions which defined the 1950s probably do not have much explanatory power in evaluating financial markets today.

To put this into perspective, here is an estimate of how many data points are available to investors operating in the stock market: The MSCI All Country World index includes a total of 2,495 constituents which are a good representation of the investment opportunity set of a global equity investor. Prices of these companies are

available on a daily basis which sums to roughly 250 prices a year. In addition, companies typically release new information on their income statements and balance sheets on a quarterly basis. Overall, there are around 300 data points every year for every company in the index universe. Over twenty years this adds up to around 15 million data points⁶ for the entire investment opportunity set.⁷

This figure is dwarfed by the amount of data required to train a deep neural network for simple categorization tasks such as distinguishing pictures of apples from those of oranges which is within the range of several hundred million data points.⁸ Furthermore, for such a task, a neural network relies on about 10 million parameters defining the function in each node in order to accurately determine the desired output. This means that for the correct parametrization of a neural network, the amount of data points available must exceed the number of parameters by far. Otherwise the training data set is too narrow and learnings could easily be based on erroneous inferences on the relevance of the sample data for the assigned task. So, it seems that in financial markets due to data shortage, the application of deep neural networks for predicting the future is very challenging. However, before jumping to conclusions, we explore the concept of predictability in financial markets in the next chapter.

⁶ The calculation $2,495 \times 300 \times 20$ results in 14,970,000.

⁷ Before you object: Using intraday data to increase the data volume will not help. Starting with Robert Merton's seminal study "On estimating the expected return on the market: An exploratory investigation" (Journal of Financial Economics, 1980), financial markets research has proven that the quality of point estimates can only be improved by increasing the length of the time period and not the frequency of the data.

⁸ Typically, an algorithm able to perform a simple categorization task consists of a neural network of 10 layers which is trained with a sample of 25,000 pictures. Assuming color pictures with 64 pixels in width and in height, we get 12,288 independent pieces of information per picture or around 300 million data points to train the network. A simple deep network has approximately ten million parameters to determine. See: sub-subroutine. (2018). Cats and dogs and convolutional neural networks. [online] Available at: <http://www.subsubroutine.com/sub-subroutine/2016/9/30/cats-and-dogs-and-convolutional-neural-networks> [Accessed 10 Apr. 2018].

Unlocking value with artificial intelligence in financial markets

Artificial intelligence joins the financial industry's chase for predictability

Considerable amounts of money can be made from predicting the market. This is a widespread hypothesis that has inspired active managers as well as academic researchers alike who have set out to develop models and tools that are able to describe and harness asset returns. And indeed, asset managers must possess some degree of predictive power in order to devise market-beating trading strategies. This is a challenge that all asset managers face be they conventional active ones or traditional quant managers or AI algorithms.

Predictability in financial markets is an elusive concept and much ink and money has been spent in an effort to define and capture it. Especially over short-term horizons, market predictability seems to be a challenge as asset returns tend to be highly volatile in the short run. Weekly intervals seem to be a little more promising since some evidence from academics and practitioners states that returns are predictable to a certain extent on a weekly basis.⁹ Over longer time spans the degree of market predictability rises in a sense that after periods of rising asset prices, the market tends to fall and vice versa.

Knowing this is not enough, however, to make money from the market which requires a more granular approach to analyzing expected market movements.

Different angles can be taken when assessing the market's degree of predictability. One is to ask how fast the market digests publicly available information. The answer to this question is tied to the Efficient Market Hypothesis that describes the degree of information efficiency in financial markets. Information efficiency is measured by the speed with which a market processes information that is, however, inversely related to its degree of predictability. If there is a lag between the release of new information and its reflection in asset prices, mispricings and, consequently, opportunities for investors arise. However, a high degree of return predictability by virtue of information advantages is at odds with competitive and efficient markets. It is, in fact, expected that competition among market participants for information advantages quickly drives out any predictable aspects to asset prices that might arise from information inefficiencies.

Another approach is to look at market participants' attitudes towards risk and their expectation to get compensated more for carrying risky assets rather than less

⁹ Lo, A. and MacKinlay, A. (1999). A non-random walk down Wall Street. Princeton, N.J.: Princeton University Press.

risky ones. There is an abundance of theories that link risk to asset prices. Today, we know that asset prices correlate with the business cycle. A negative economic outlook comes with low asset prices, high risk premia and, therefore, high return expectations. The opposite is true when an economic upswing is expected. Return predictability derived from risk premia rests on assumptions on compensation for risk which allows for the predictability of asset prices, even in efficient markets. However, just because risk premia are predictable, investing successfully is not a trivial undertaking.

Asset managers and researchers have often seen their assumptions overturned by unexpected market movements that contradict widely accepted asset pricing models and their predictive aspirations. Moments like these are naturally grist to the mill for those skeptical academics and practitioners who are strong advocates of the hypothesis that asset prices are in fact unpredictable and random. So, despite it being so sought after by market participants, predictability in financial markets is, after almost half a century of research, still difficult to pin down.

Why are we so convinced then that financial markets can be predicted? Much of what we consider evidence on the market's predictability comes from the practice of backtesting. Commonly, investment strategies are tested by using a sample of past returns with the goal of simulating future performance. Although widely accepted among practitioners and academics, this approach is fraught with pitfalls that might overstate the future profitability of a given investment strategy. Overfitting occurs when an analysis corresponds too closely to a particular data set and fails in out-of-sample testing. Basically, if you have fallen prey to overfitting, you think a strategy is very likely to be successful in the future, when it actually only succeeds in a particular scenario of the past. However, this is a potential dilemma that is unlikely to be resolved since all investment practitioners are inevitably chained to the past which might give rise to illusory perceptions of predictability in financial markets.

In particular, AI driven investment strategies are prone to overfitting. This is because their underlying algorithms are made up of a large number of parameters. On the one hand, this allows for a high degree of model flexibility which increases the chance that the algorithm, once trained with sample data, will produce accurate results when applied to new data sets. On the other hand, this feature makes the algorithm more susceptible to overanalyzing data and data noise. This can result in the detection of data patterns that only work in a specific data sample and that cannot be transferred to new data sets. So, investors need to be cautious when judging the value of AI strategies. There will always be a promising AI algorithm that beats the market in backtesting.

In addition, AI strategies often lack the one remedy that has proven to be efficient against the pitfalls of backtesting: a sound theoretical foundation in financial markets theory. Strategies firmly rooted in financial market theory

“AI driven investment strategies are prone to overfitting.”

are backed by a robust rationale and are less dependent on assumptions made from data analysis. While this does not bode well for AI strategies, it does not mean that they have no explanatory power in our markets. Instead, what this circumstance demands is an in-depth understanding of how and when AI can make a credible contribution to the practice of investing.

How can AI add value?

In order to obtain a high degree of predictive power, AI, as any other market participant, must be able to accurately anticipate future movements of the factors that drive asset prices. Take a company's stock price for example. A stock price is a reflection of three components that can influence its movements. These are news impacting corporate cash flows, changes in the rate of return required by investors to hold the stock (also called the discount rate which consists of the risk-free rate of return and an equity risk premium) and behavioral biases of market participants. An example for each will make those return drivers more tangible:

- News on corporate cash flows: In November 2017, the pharma giant Roche Holding AG was successful in testing two of its cancer drugs in late-stage patient drug trials. Shortly after the announcement, the company's stock surged by 6 percent as market participants were eager to get a piece of the cake with the objective to participate in Roche's higher future earnings potential.
- Changes in the rate of return: In December 2017, the U.S. Federal Reserve announced a rise of the federal fund target rate, from 1.25 to 1.50 percent. Central bank reference rates are the major determinant of the discount rate used to discount expected future corporate cash flows back to the present. So, equity markets may react strongly to these kinds of changes. At the time of the announcement, the stock market reaction remained muted since the rate hike was in line with market expectations that had anticipated a rise to 1.49 percent. Typically, prior to the announcement, there is a high degree of speculation going on

→ continued on page 17

How AI became the game master of chess

Artificial intelligence snatched the chess crown from human hands a long time ago. In 1997, IBM's supercomputer Deep Blue defeated world chess champion Garry Kasparov under tournament conditions. It was the first time that a machine beat the best human player at the legendary strategic board game. Already in 1769, the idea of a machine beating a human at chess fascinated human imagination when the Hungarian nobleman Wolfgang von Kempelen built the Mechanical Turk which was a fake chess-playing machine and one of the most elaborate hoaxes of humankind.

Today, beating a human at chess is easy for a computer and pitching computers against humans has become a dull affair. Computer-only chess leagues have been the result which, until fairly recently, were headed by the open-source chess engine Stockfish. In December 2017, however, it was beaten by DeepMind's AlphaZero in a 100-game tournament. AlphaZero is an intelligent game-playing machine which was derived from its better known sibling AlphaGo that has repeatedly defeated the world's best players of the ancient Chinese board game Go. Within just four hours AlphaZero taught itself how to play chess with no human input and succeeded in winning 90 games against the world leading chess-playing engine.

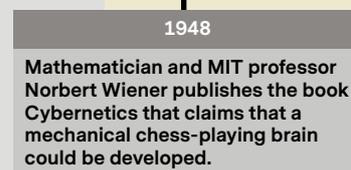
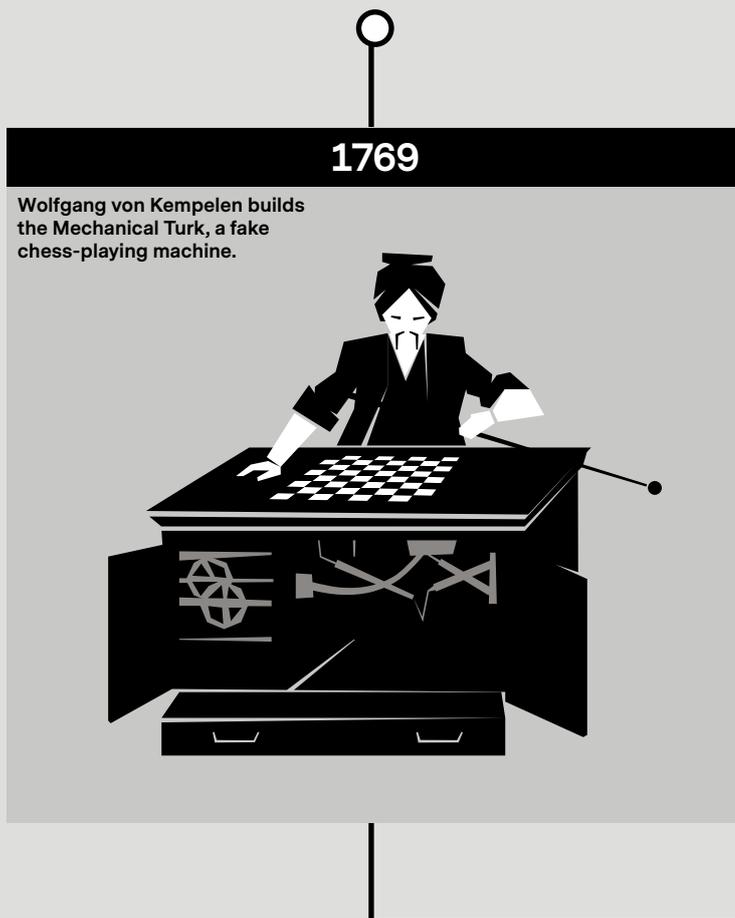
Chess is a complex game

The difficulty in chess comes from the challenge to identify the optimal strategy to win the game from a multiplicity of possible moves in a deterministic game environment of perfect information. This means that both players are perfectly informed on all game states at all times. In addition, there are no random elements to the game so all player actions have a clearly defined outcome. It is generally accepted that solving chess by finding the ultimate strategy that always wins, is not possible due to the sheer number of possible games and positions that can be taken by a player in the course of a game.

The American mathematician Claude Shannon evaluated the game tree complexity of chess.* He calculated the number of possible games based on the assumption that each game consists of an average of 40 moves and that for each move a player decides between 30 possible moves. What has become known as the Shannon number amounts to 10^{120} possible games in chess – a number which exceeds even the current estimate of observable atoms in the universe. Good human chess players tend to look ahead five “plies” which is generally required to play at grandmaster-level. “Ply” is chess speak for a move by an individual player. Five might seem a small number. Don't let yourself be duped though. After five moves, almost 5 million board positions are possible.

While humans can rely on intuition to choose their moves, computers must follow a strictly systematic approach. They look at chess as a game tree and are programmed to find the optimal winning strategy in any given environment as fast as possible. Solving chess by considering all possible moves, however, is a task at which even computers fail. A computer that can evaluate 1,000,000 positions per second would take more than 30 years to examine ten plies.

* A game tree is a graphical representation of a game consisting of sequential moves. It starts with the initial position of a game and shows all possible moves from that position resulting in a new position from which flows a multitude of further moves.



Why AI is so good at chess

Chess represents a well-defined, rules-based game environment that lends itself well to illustrating how artificial intelligence has gained the upper hand over human brain power in one of the most complex and intellectually demanding games of our times. The rapid increase in computer processing power which has been driven by Moore's Law over the past decades, has teamed up with methodological improvements in underlying game-solving algorithms to propel AI's ability to master chess. In 1983, Belle was the first chess machine to achieve master-level chess by pure computational processing power thanks to its ability to evaluate 100,000 positions per second. However, it was only improvements in quantitative game-solving algorithms that made it possible for Deep Blue to evaluate up to 200,000,000 positions per second and win over the best human player of all time.*

* Research.ibm.com. (2018). IBM Research | Deep Blue | Overview. [online] Available at: <https://www.research.ibm.com/deepblue/meet/html/d.2.shtml> [Accessed 8 Mar. 2018].

1950
Claude Shannon publishes *Programming a Computer for Playing Chess*. He demonstrates that trying to solve chess by brute force is pointless. The Shannon number becomes known.

1951
Alan Turing is first to specify a chess playing algorithm dubbed *Turochamp*. Back then no machine existed that could implement the program.

1956
John McCarthy invents an algorithm called *alpha-beta search* that helps computers get better at chess.

1962
The first chess program to play convincingly, *Kotok-McCarthy*, is published at MIT.

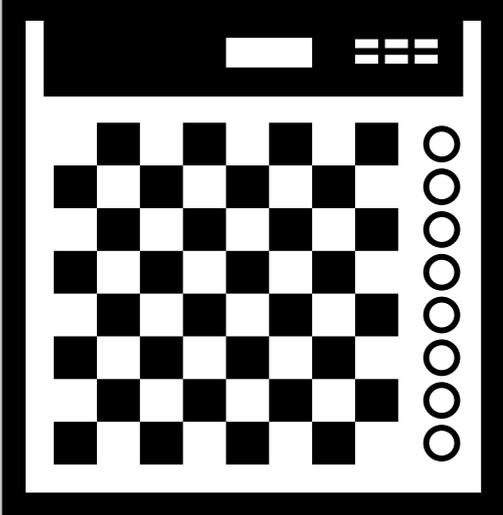
1966/67
The first chess match pitting computer against computer is played. *Kotok-McCarthy* loses against the program developed by the Moscow Institute for Theoretical and Experimental Physics.

2017
AI Elmo wins at the 27th Annual World Computer Shogi Championship.

1997
Deep Blue wins a six-game match against Garry Kasparov.

1996
Deep Blue loses a six-game match against Garry Kasparov.

1977
The first chess playing micro-computers, *Chess Challenger* and *Boris*, are released.



The illustration shows a chessboard with alternating black and white squares. To the right of the board is a vertical column of eight white circular pieces. Above the board, there are two small rectangular boxes, one containing a white square and the other containing three horizontal lines, representing a chess clock display.

2017

AlphaGo defeats the world's top Go player Ke Jie.

2017

Alpha Zero beats Stockfish in a 100-game tournament.

2017

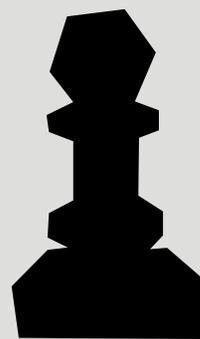
Libratus wins at Poker against 4 human top players.

2018

Alpha Zero teaches itself to play three different board games (chess, Go, and shogi) in just three days.

2019

DeepMind's AI AlphaStar beats human professionals ten games in a row in the real-time strategy video game "Star Craft II".



Computers step up their game

To speed things up, chess computer programs have resorted to quantitative game-solving algorithms that have seen considerable improvements over time thanks to progress in technical engineering. Traditional methods include plain search algorithms that exhaustively explore all move combinations along the game tree in order to select the best one. As this approach takes time, it was enhanced by a method called alpha-beta-pruning which allows to eliminate less promising move combinations more quickly. This frees up time to evaluate more useful game tree branches. Deep Blue and Stockfish relied on search algorithms and their improved variants. They succeeded at chess by brute force. They evaluated as many game tree paths as possible at the highest possible speed. However, even in their optimized form, search algorithms must proceed in a strictly sequential fashion along the game tree. This means that computers can analyze the game tree only to a certain depth in a given amount of time since the amount of move combinations grows exponentially with increasing game tree depth.

Self-reinforcement learning algorithms beat everything else

This is why they were swiftly outsmarted by AlphaZero which capitalizes on benefits from new developments in quantitative game-solving algorithms. Powered by self-reinforcement learning AlphaZero no longer treats chess as a search but as learning problem. In the context of games, self-reinforcement learning algorithms have no other knowledge than the rules of the game and they learn how to play well by repeatedly playing against themselves. They are able to educate themselves since the outcome of the game acts as a feedback mechanism by providing punishment (failure) or reward (success) for chosen actions. Basically, the machine learns which moves are good and which ones are bad in any given game situation. Good board positions are remembered while bad ones are discarded.

Relying on this method, AlphaZero taught itself to play chess in a few hours. It was programmed to play randomly chosen game tree branches with the goal of learning millions of different chess moves, board positions and their outcomes. This way it assembled a cluster of associations which, under real tournament conditions, helped to assess previously unknown game situations and deduct the most promising moves in a particular environment. Self-reinforcement learning algorithms do not consider chess as a computational

challenge which requires the exhaustive evaluation of all possible moves. Rather, they treat it as a learning environment which demands the maximization of reward at the end of the game.^{*} This way, self-reinforcement learning algorithms can evaluate game trees faster, smarter and in greater depths than their predecessors that rely on traditional game-solving methods.

AI's superiority in chess is no magic

AI's supremacy in chess over humans purely results from superior calculation abilities and human-engineered game-solving algorithm improvements that are systematically applied to a game tree. So, AI-driven computer power in chess is no magic; it is down to computational speed and state-of-the-art technical engineering. Over the past few years computers have started to tackle more complex deterministic games such as Go and the Japanese version of chess, shogi – and they have succeeded.

More recently, even non-deterministic games of imperfect information such as poker or bridge, have come within AI's reach. In 2017, an AI-driven computer called Libratus won out over four of the world's best poker players at a Pittsburgh casino. Should AI fully gain mastery of data processing in imperfect information environments, financial markets could become a promising area of application for AI algorithms. However, the answer to the question if AI will score high at the financial markets game, will be left to the human engineers crafting its intelligent algorithms.

* Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm (2017). [ebook] London: Deep Mind. Available at: <https://arxiv.org/pdf/1712.01815.pdf>,%20Minimax%20search%20and%20reinforcement%20learning%20for%20adversarial%20tetris [Accessed 8 Mar. 2018].



driving asset prices in the expected direction. When the change is finally announced and market expectations are not met, prices react within seconds.

- Behavioral biases: In only two months, in November and December of 1999, the NASDAQ returned more than 40 percent pushing the total return of the year above the 100-percent mark. This instance is a symptom of the dotcom bubble which is a prime example for herding. Herding is a behavioral bias where a large number of market participants gravitate to the same kind of investments merely based on the fact that those around them are investing in these securities.

The power of AI in predicting these fundamental return drivers resides with its ability to assess information faster and deeper than humans. Based on the law of active management AI has the potential to generate value in two ways: more and better bets. This results in a more comprehensive data analysis and capture. AI can process larger amounts of data at higher speed which enables it to detect data patterns or linkages that would otherwise go unnoticed. This ability is a coveted skill when capitalizing on what we call the big data revolution. The rapid

“AI has the potential to generate value in investing in two ways: more and better bets.”

increase in the amount of data available together with the strong expansion of data storage capacities has opened up new playground to those who can digest vast amounts of unstructured data in an efficient and intelligent manner. Examples of how AI can improve predictions on the above mentioned three fundamental stock price drivers is presented below:

- Natural language processing by intelligent algorithms can extract value from unstructured data on a company that comes in the form of annual or quarterly reports, news reports, social media posts, emails and the like. By making inferences from language, machines can discern information on a specific company which allows it to make projections on the company’s revenue and cash flow streams. As long as this information is not fully priced in yet, profits can be made.¹⁰ However, in order to get a head-start, a large data set and the appropriate toolbox to quickly and systematically process the data are required.

- Much brain power goes into evaluating the current and future state of the economy. Economists and investors relentlessly screen and scrutinize a wide range of economic indicators for this purpose. The challenge they face is that most of the economic indicators such as inflation rates, GDP and employment data are published infrequently with a substantial time lag and low statistical precision. This is why forecasting the macroeconomic conditions has taken the back seat to what is called the practice of “nowcasting”. Thanks to their capacity for large data volumes, intelligent algorithms could substantially increase the number and types of indicators monitored as well as the length of the time series examined. Applied to the field of macroeconomic analysis, AI could improve the quality and precision of assessing the current state of the economy which helps to gauge future interest rate moves by central banks, current risk premia and the overall discount rate more accurately. By extension, precise forecasts of central bank rates improve equity price projections since the discount rate applied is more exact.¹¹

- Forecasting investor sentiment could be lucrative. This is because research has shown that sentiment-induced price drops are usually followed by an upswing in prices. Since investor sentiment is reflected in online and print media, having an intelligent machine screen vast amounts of public text sources could detect emerging investors’ sentiment already in its beginnings. Media monitoring performed by humans can only capture bits and pieces of a larger picture and might catch on to emerging investor sentiment only too late – only after it has already moved the market. AI algorithms could be put to use in order to collect immense amounts of sparse text data, aggregate, evaluate and learn from it with the goal of deriving actionable investment insights before investor sentiment impacts the market.

From the above, it appears that AI is well equipped for the task of obtaining a high degree of predictive power on what drives stock prices. Placing more and better bets on markets makes for improved information capture in large data pools. Opening up new ways of utilizing big data, AI has the potential to convert unstructured data into valuable inputs for profitable trading strategies. However, certain conditions on a market’s degree of information efficiency must be met for AI to have a bearing on investing successfully.

¹⁰ Tetlock, P., Saar-Tsechansky, M. and Macskassy, S. (2007). More than Words: Quantifying Language to Measure Firms’ Fundamentals. SSRN Electronic Journal.

¹¹ Bok, B., Caratelli, D., Giannone, D., Sbordone, A. and Tambalotti, A. (2017). Macroeconomic Nowcasting and Forecasting with Big Data. [ebook] Federal Reserve Bank of New York. Available at: https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr830.pdf [Accessed 6 Mar. 2018] and Ng, S. (2016). Opportunities and Challenges: Lessons from Analyzing Terabytes of Scanner Data. SSRN Electronic Journal.

When can AI add value?

The Efficient Market Hypothesis describes three different types of market structures that are characterized by varying degrees of informational content. The more information a market reflects in asset prices, the more efficient a market is. More precisely, the information contained in asset prices may range from “past information only” via “public information only” to a full reflection of “public and private information” depending on the prevailing market structure.

A useful tool in assessing a market’s efficiency is the analysis of how quickly prices mirror new information. If all available information was instantaneously reflected in asset prices, there would be no advantage to be gained by anyone. For a market to be efficient in information processing, it must have certain qualities. Market participants must be able to trade freely and at low cost to react to news quickly. In addition, the majority of market participants must act in a rational manner in order to be able to assess the impact of news on corporate cash flows and act accordingly. Given these conditions market participants are able to swiftly convert information into actionable investment insights, trade on them and move asset prices until they reflect all relevant information.

It looks like a market has quite a few prerequisites to meet in order to be considered informationally efficient. The above criteria will not always be given at all times. Assessing AI strategies in a scenario analysis will provide insights as to its potential for impact.

1. Strong form of market efficiency

In strong form efficient markets, all public and private information is instantaneously reflected in asset prices. Since there is no time lag between the release of information and its incorporation by the market there is no informational advantage to be gained by any investment or trading strategy. There is no such thing as private information and the market cannot be outperformed. There is simply no information to bank on. Costly information gathering would in fact result in negative returns.¹³ As any other investor, AI would be left with the only option of passively tracking the market. However, AI is unlikely to encounter this form of market efficiency since empirical evidence does not provide any support for the existence of such markets. Hence, as Eugene Fama rightly puts it, since “the extreme version of the market efficiency hypothesis is surely false,”¹⁴ an AI algorithm passively investing in the market is an unlikely scenario.

2. Semi-strong form of market efficiency

Under this form of market efficiency, the market adjusts quickly to absorb all publicly available information. Only if market participants possess non-public information, can they achieve abnormal returns. As a consequence, research, which is able to collect information unavailable to the public, has value. However, from an investor's perspective, value is only added if research costs do not erode the additional gain in performance. This is a challenge that is difficult to resolve given that it takes about 30 to 50 research analysts to cover the entire 2,495 companies that make up the MSCI All Country Index and to form a solid opinion on them. The application of AI algorithms to unstructured text data on the same investment universe could be a cost-efficient way to monetize the data by turning it into valuable information that is not reflected in public prices. Should AI prove to be a good researcher, analysts will need to reinvent themselves in order to harness the power of this new tool.

3. Weak form of market efficiency

Advocators of weak form efficient markets state that future asset prices are random and not influenced by past information. They assume that past price movements have no bearing on future asset prices. New information changes prices and since new information is random, prices changes are random. As a consequence, this view rejects technical analysis as a valid tool to uncover price development patterns that are likely to be repeated in the future. For example, the belief is that momentum in stock prices does not exist.

However, the hypothesis of weak form efficient markets is not well supported by academics and market participants.¹⁵ In fact, it runs counter to a large group of investment practitioners that have successfully pursued systematic trend following strategies since the 1970s. Commodity Trading Advisors (CTAs) are technical traders who buy upward trending assets and sell downward trending assets betting on these trends to persist. For example, in a down market, they count on prices being driven mainly by cascading effects since investors face adverse conditions such as risk budget restrictions and margin calls, which is expected to result in strings of neg-

ative returns that can be exploited.¹⁶ Sophisticated CTAs are fully automated investment managers applying a wide range of methods, often borrowing directly from the AI toolkit. Considering AI's strength in pattern recognition, CTAs are likely to increase their reliance on AI technology to improve their quantitative trading strategies.

“Data is the fuel to AI's engines.”

In sum, AI algorithms possess superior information processing and data pattern detection capabilities which empower them to formulate accurate predictions on asset prices. Any prediction on financial markets is only useful, however, if there is room for it to have an impact in the market. This means that in strong form efficient markets, AI strategies would have no advantage over any other asset manager and blend in with mainstream passive investing. In contrast, semi-strong and weak form efficient markets do provide the necessary leeway for AI to exploit the realms of non-public information that exist under these market structures – however small and fleeting they may be. Under these conditions, AI can pursue unconventional ways of assembling public information from known as well as unexplored data sources. The goal is to form new data clusters representing pieces of information that are not yet reflected in asset prices. This way, AI is able to harness the power of non-public information by converting large data sets into actionable investment insights in a cost-efficient way.

While this seems like a veritable treasure trove for AI driven asset managers, unlocking the vast domain of pattern recognition in financial analysis by means of AI algorithms could release a true gold rush into their clients' portfolios. Data is the fuel to AI's engines, so if past data is believed to harbor valuable information for the future, AI could shift into a higher gear when applied to technical analysis. Limits to AI's potential in this area are set by how valid we assume the weak form of efficient markets to be.

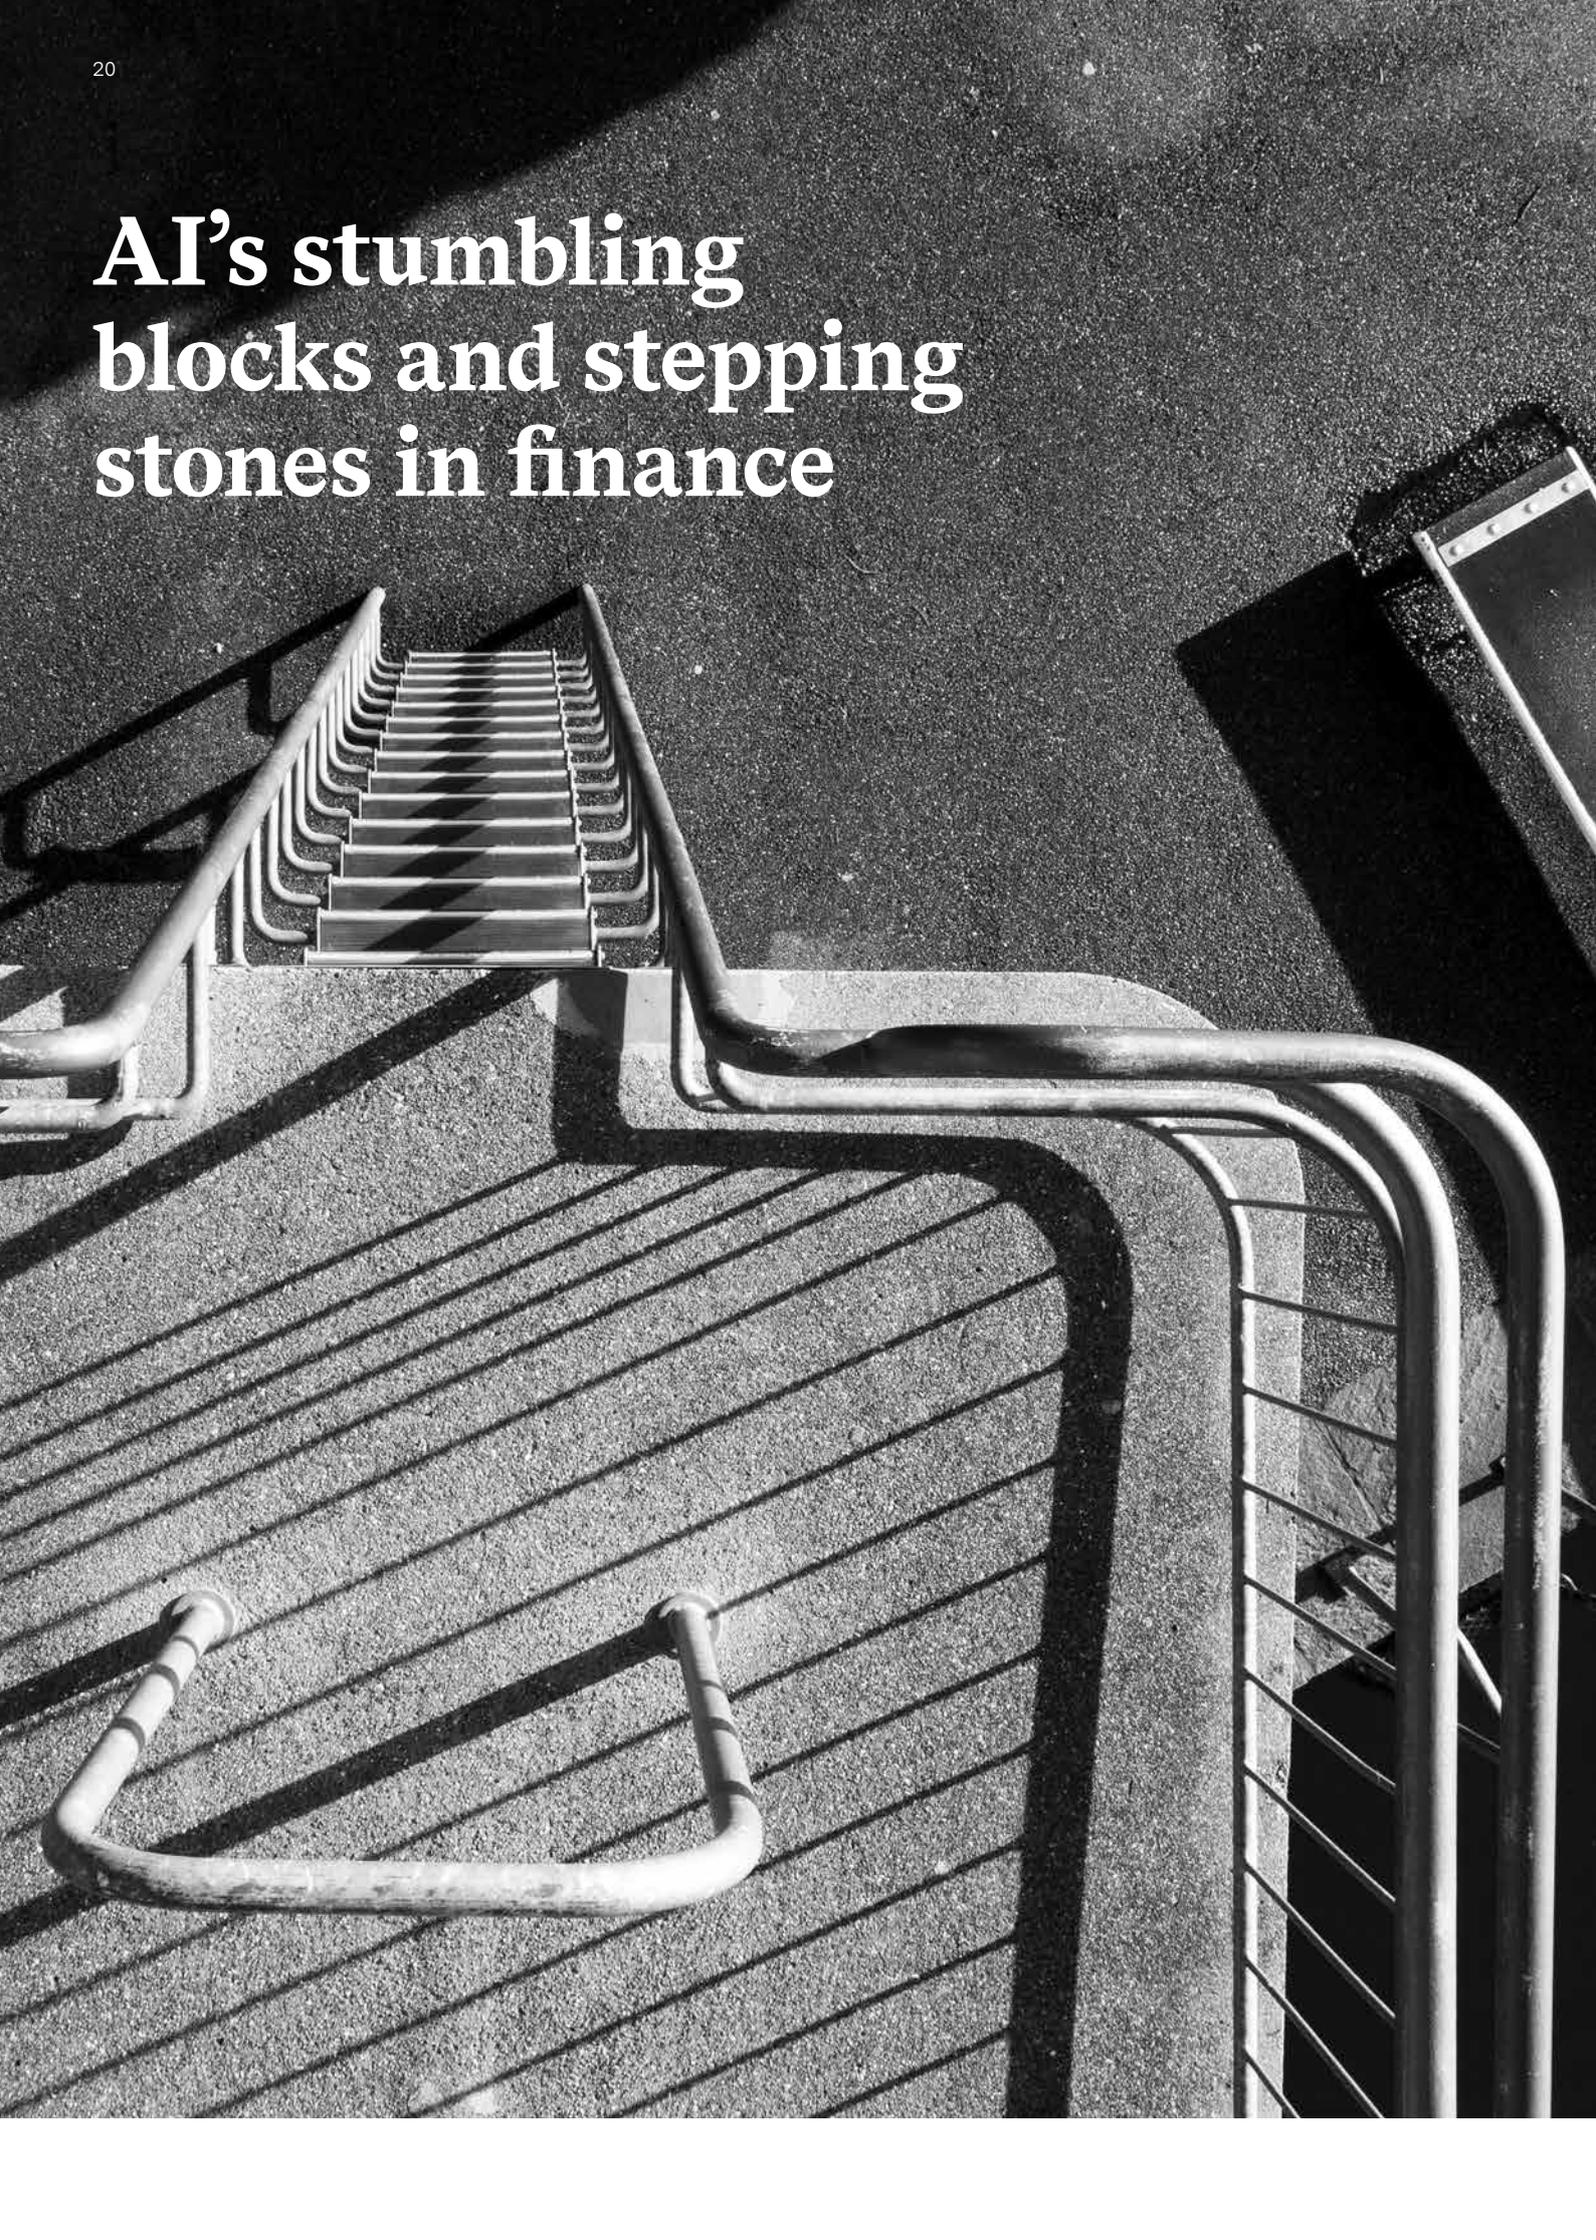
¹³ Grossman, S. and Stiglitz, J. (1980). On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 70(3), pp. 393–408.

¹⁴ Fama, E. (1991). Efficient Capital Markets: II. *The Journal of Finance*, 46(5), p. 1575.

¹⁵ Jegadeesh, N. and Titman, S. (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance*, 48(1).

¹⁶ A more theoretical explanation of the cascading effect can be found here: An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies, *The Journal of Business*. Grossman, S. (1988). An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies. *The Journal of Business*, 61(3), p. 275.

AI's stumbling blocks and stepping stones in finance



New techniques in the artificial intelligence toolkit harbor potential pitfalls

Accurate predictions of financial market movements are the holy grail of investing. All financial analysis methodologies, be they traditional statistical techniques or AI algorithms, share predictive accuracy as their ultimate goal. Common AI methods such as regression analysis as well as clustering and classification methodologies have been used by quant managers around the globe for decades. AI changed face with the emergence of big data and the exponential rise in computer processing power which enabled advances in machine learning. With it came a key feature that separates new AI methods from traditional statistical ones that are most commonly used in econometrics for investment purposes. This feature could prove to be a trap for AI driven quants.

The key difference between the two approaches to financial analysis lies in the way of how predictions are derived. Traditional methods in econometrics always strive to establish a causal link between the prediction made and the reasons for it. They commonly trace back their claim to theoretical foundations or economic conditions in order to establish credibility. In other words, not only the highest quality point estimate of a variable is the aim but also its best explanation. In contrast, modern AI methods are purely outcome-driven. They solely focus on the best point estimate. No explanation is given, which is why investment strategies that bet on intelligent algorithms run the risk of turning into incomprehensible black boxes that arouse suspicion in investors who increasingly demand transparency from those organizations that they entrust their capital to.

The black box problem

Due to their outcome-driven nature, many AI based decision-making processes remain obscure. Even simple recommender systems, such as online streaming systems suggesting movies according to their users' profiles, are not fully transparent on how they work. This shortcoming affects deep learning algorithms in particular that, due to their depth and complexity, are a quagmire of mathematical functions and calculations that are not even comprehensible to their creators, much less to their users. Even though these networks are built by humans, training is

done by the models themselves with impressively accurate outputs. However, the process of how these outputs are derived cannot fully be explained. Even what makes a network ultimately recognize a face on a picture is still hard to say. This does not inspire confidence. Just as much as users want to understand how a self-driving car makes its decision on when to brake, investors will want to know why an algorithm chooses to invest in equities rather than in bonds at a certain moment in time.

“Transparency can be defined as the disclosure of requested information at the required time.”

Lack of transparency might prove to be debilitating for AI-based decision-making processes in certain industries. One of the reasons is the increased regulatory activism of EU officials who, in May 2018, introduced the General Data Protection Regulation which marks a watershed in data-privacy regulation in Europe. The regulation applies to any organization processing data from EU citizens. Article 22 of the regulation contains a provision that requires companies to provide users with an explanation of how an automated decision was reached by an intelligent algorithm. Regulators go even further by enshrining a fundamental right to explanation in the regulation whereby users are entitled to ask for an explanation from an algorithm for decisions that significantly affect them. This puts many algorithms currently in use on the spot since their developers will be at a loss for words when attempting to shed light on their inner workings. Algorithms may have to be overhauled and new standards for their development may have to be set while others might have to be shut down entirely in order to keep up with the current trend towards increased transparency in the investment management industry.¹⁷

The importance of interpretability of automated decision making processes is particularly pressing in finance when investment decisions are made. All lasting relationships between financial institutions and their clients are based on trust which is built either on clients understanding the investment processes behind their portfolio or clients

¹⁷ For more information, please refer to Vescore's white paper "Why quant?"

having faith in the abilities of their asset managers. The enabler of each of these two conditions is transparency which can be defined as the disclosure of requested information at the required time.

Information can be made available on the investment portfolio and the investment process. While the investment portfolio can be explained by simply revealing the current asset allocation down to single instrument positions, divulging the intricacies of the investment process is more complicated. Besides disclosing who makes what kind of decisions, it involves explaining the reasons behind allocation changes. This is where many AI investment algorithms fail due to their inability to translate their algorithmic design into language that a non-expert user can understand. In contrast, a human investment manager will be able to provide an explanation for his investment decisions which, despite the fact that some parts of it may remain rather instinctual and not fully comprehensible to everyone, is more likely to instill confidence in investors than a black box algorithm whose reasoning is inscrutable. So, is the main advantage that the human manager has over AI the ability to tell a story?

Our ability to tell stories differentiates us from machines

While the increasingly sophisticated nature of tasks performed by AI seems to indicate that the gap between humans and machines is closing, narrative intelligence might well be what sets human minds apart from their artificial counterparts. Patrick Henry Winston's Strong Story Hypothesis confirms this view as it suggests that the ability to "tell, understand and recombine stories"¹⁸ is what is unique to humans. Storytelling enables humans to relate to others and to their environment. Stories are how we experience our world. It is not only the imaginative power to create a story from a string of words, that conjures up mental images and evokes emotions, which is unique to the human mind. But also the ability to understand stories, chain events together and unchain them to manipulate, reflect on and explain stories is what makes human intelligence special. So as long as computerized machines are not able to craft and understand stories and respond to them in an empathic manner, they will continue to lag behind the human mind.

The white box solution

Storytelling as a communications tool is at the core of the financial services industry. Financial institutions lack physical products and, therefore, they depend more heavily on the tools of communication in order to sell their concepts and build meaningful client relationships. Bridging the gap between abstract investment processes and clients seeking hands-on and explainable returns is of the utmost importance for investment managers. This applies not only during periods of underperformance when investor queries abound but also when performance is good. Investors will not have faith if they perceive strong outperformance as a potential hoax.

“Investor trust is the prerequisite for success in the financial services industry.”

Transparency is achieved by means of communication that puts the intellectual products and processes of the industry into words that clients can comprehend. So, if AI based investment strategies want to be recognized as viable alternatives by investors to other quants or active managers, they must establish full transparency on their investment decisions and how they are derived. This is why it is imperative that AI go beyond making investment decisions and build dynamic communication interfaces with investors with the purpose of translating its web of mathematical functions into investor-friendly intelligence. If an algorithm acquires the ability to explain the reasons for its investment decisions via an artificial interactive agent, such as an intelligent chatbot, it lays the foundation for investor trust which is the prerequisite for success in the financial services industry. AI powered communication agents are a potent approach to transforming a black box algorithm into a white box investment engine. The advantage of artificial interactive agents is that they are free from bias and personal perceptions that human intermediaries are prone to. Messages that are fed from one computerized network to another do not get distorted and preserve their original meaning.

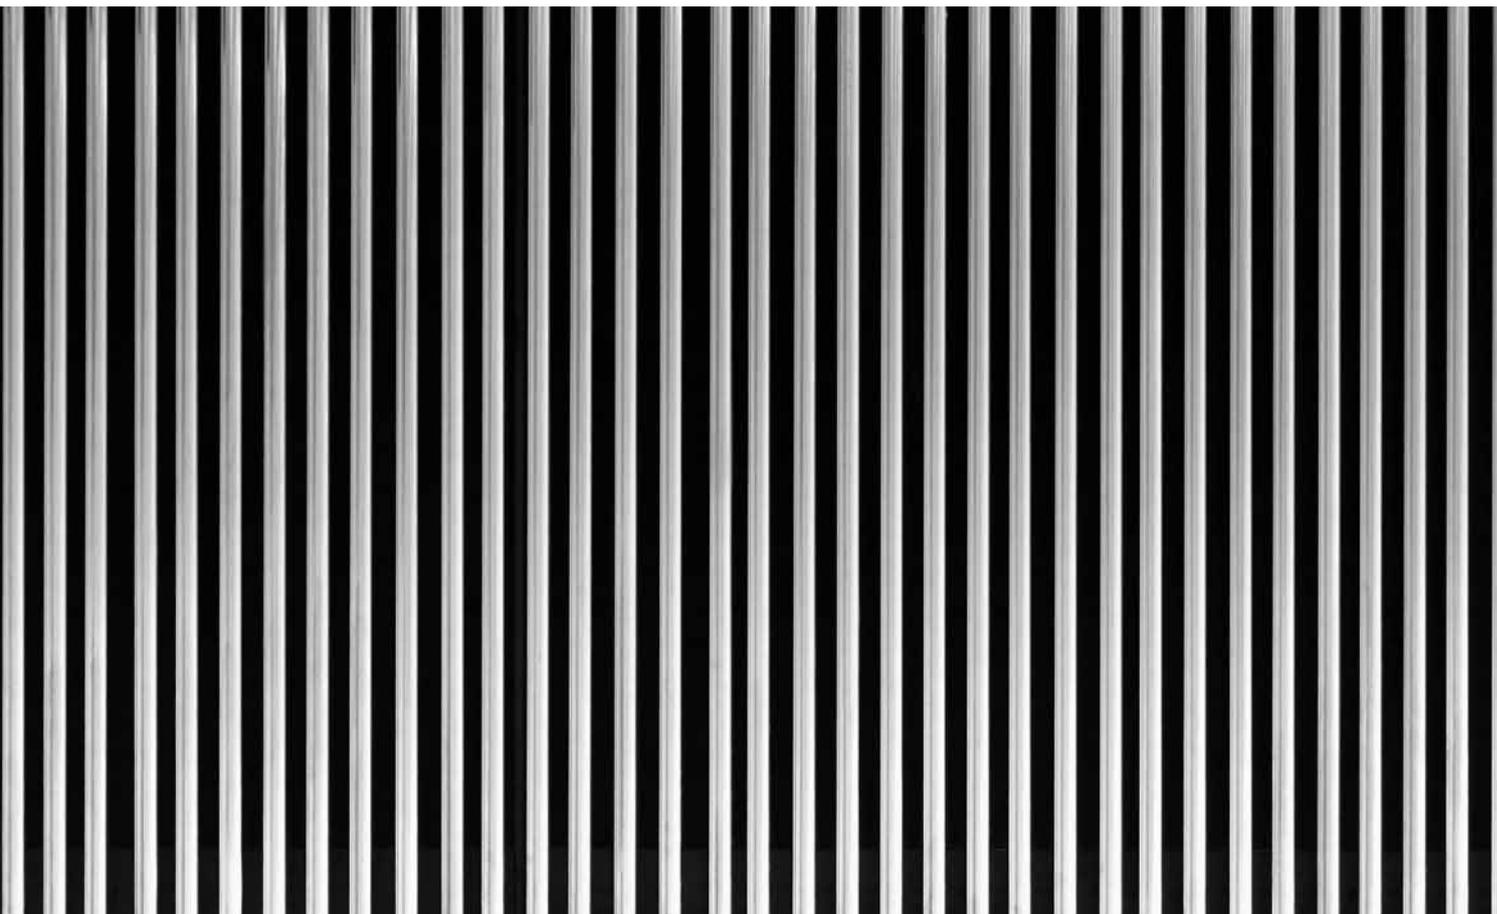
¹⁸ Winston, P. (2018). The Strong Story Hypothesis and the Directed Perception Hypothesis. [ebook] Massachusetts Institute of Technology. Available at: <https://www.aai.org/ocs/index.php/FSS/FSS11/paper/viewFile/4125/4534> [Accessed 13 Jan. 2018].

“AI powered communication agents are a potent tool in transforming black box algorithms into white box investment engines.”

Once AI is able to spell out the details of an algorithmically motivated investment decision straight from the market condition that triggered a signal to the neural network through to the output result, it becomes accountable for its decisions and is able to satisfy investors and regulators that both are adamant on the subject of transparency. Ultimately, an AI algorithm embedded in a fund that, at the mere click of a button, interacts with the client and tells the story of its investment decision without any human intermediary, will become the epitome of transparency in investing.

The future of artificial intelligence in investing is bright

Investment funds telling their stories could be just the prelude to new practices in the investment industry. With its prowess in deriving insights from data and processing capacities, AI is poised to invade many areas of investing that go beyond exploiting data patterns and generating investment insights. AI could revolutionize not only the way single portfolios are managed but also how multi-manager portfolios are maintained. In the future, we could see self-organizing portfolios that, based on interacting algorithms, could assemble various funds in an efficient portfolio management process that avoids inefficiencies and redundancies while at the same time adhering to restrictions defined by the client in the areas of asset allocation and risk management. The possibilities for AI to take a foothold in investing are manifold. Yet, ultimately, the future of AI in investing depends on its ability to make money for investors in a transparent and intelligible way.



Reference List

- A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play (2018). David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, Demis Hassabis. *Science* 07 Dec 2018. Vol. 362, Issue 6419, pp. 1140 – 1144. DOI: 10.1126/science.aar6404
- AlphaStar (2019). Mastering the Real-Time Strategy Game StarCraft II <https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii>
- Bok, B., Caratelli, D., Giannone, D., Sbordone, A. and Tambalotti, A. (2017). *Macroeconomic Nowcasting and Forecasting with Big Data*. [ebook] Federal Reserve Bank of New York. Available at: https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr830.pdf [Accessed 6 Mar. 2018].
- Boughin, J., Manyika, J. and Woetzel, J. (2017). *A future that works: automation, employment and productivity*. [ebook] McKinsey Global Institute. Available at: https://www.mckinsey.com/mwg-internal/de5fs23hu73ds/progress?id=6TfTsrwVVIPeNn1ecR_ebKTI-APjXO-wNzVaSWZPq1z0,&dl [Accessed 10 Feb. 2018].
- Fama, E. (1991). Efficient Capital Markets: II. *The Journal of Finance*, 46(5), p.1575.
- Grossman, S. and Stiglitz, J. (1980). On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 70(3), pp. 393–408.
- Grossman, S. (1988). An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies. *The Journal of Business*, 61(3), p. 275.
- Jegadeesh, N. and Titman, S. (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance*, 48(1).
- Lo, A. and MacKinlay, A. (1999). *A non-random walk down Wall Street*. Princeton, N.J.: Princeton University Press.
- Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. (2017). [ebook] London: Deep Mind. Available at: <https://arxiv.org/pdf/1712.01815.pdf,%20Minimax%20search%20and%20reinforcement%20learning%20for%20adversarial%20tetris> [Accessed 8 Mar. 2018].
- Ng, S. (2016). Opportunities and Challenges: Lessons from Analyzing Terabytes of Scanner Data. *SSRN Electronic Journal*.
- Research.ibm.com. (2018). IBM Research | Deep Blue | Overview. [online] Available at: <https://www.research.ibm.com/deepblue/meet/html/d.2.shtml> [Accessed 8 Mar. 2018].
- sub-subroutine. (2018). Cats and dogs and convolutional neural networks. [online] Available at: <http://www.subsubroutine.com/sub-subroutine/2016/9/30/cats-and-dogs-and-convolutional-neural-networks> [Accessed 10 Apr. 2018].
- Tetlock, P., Saar-Tsechansky, M. and Macskassy, S. (2007). More than Words: Quantifying Language to Measure Firms' Fundamentals. *SSRN Electronic Journal*.
- Tetlock, P. (2007). Giving Content to Investor Sentiment: The Role of Media in the Stock Market. *The Journal of Finance*, 62(3), pp. 1139–1168.
- Winston, P. (2018). *The Strong Story Hypothesis and the Directed Perception Hypothesis*. [ebook] Massachusetts Institute of Technology. Available at: <https://www.aaii.org/ocs/index.php/FSS/FSS11/paper/viewFile/4125/4534> [Accessed 13 Jan. 2018].

Important legal information

This document is for information purposes only and nothing contained in this document should constitute a solicitation, or offer, or recommendation, to buy or sell any investment instruments, to effect any transactions, or to conclude any legal act of any kind whatsoever. Although Vontobel Asset Management AG ("Vontobel") believes that the information provided in this document is based on reliable sources, it cannot assume responsibility for the quality, correctness, timeliness or completeness of the information contained in this document. Except as permitted under applicable copyright laws, none of this information may be reproduced, adapted, uploaded to a third party, linked to, framed, performed in public, distributed or transmitted in any form by any process without the specific written consent of Vontobel. To the maximum extent permitted by law, Vontobel will not be liable in any way for any loss or damage suffered by you through use or access to this information, or Vontobel's failure to provide this information. Our liability for negligence, breach of contract or contravention of any law as a result of our failure to provide this information or any part of it, or for any problems with this information, which cannot be lawfully excluded, is limited, at our option and to the maximum extent permitted by law, to resupplying this information or any part of it to you, or to paying for the resupply of this information or any part of it to you. Neither this document nor any copy of it may be distributed in any jurisdiction where its distribution may be restricted by law. Persons who receive this document should make themselves aware of and adhere to any such restrictions. In particular, this document must not be distributed or handed over to US persons and must not be distributed in the USA.

Vontobel
Gotthardstrasse 43, 8022 Zurich
Switzerland

Vontobel Asset Management S.A.
2-4, rue Jean l'Aveugle, 1148 Luxembourg
Luxembourg

Vontobel Asset Management S.A.
Niederlassung München
Leopoldstrasse 8-10, 80802 Munich
Germany