General principles for the use of Artificial Intelligence in the financial sector

DeNederlandscheBank

EUROSYSTEEM

General principles for the use of Artificial Intelligence in the financial sector

Author: Joost van der Burgt

Contents

Practical information

Artificial Intelligence (AI) is increasingly being applied in the financial sector and is likely to become more important in the years to come. How the adoption of AI will come to shape the financial sector precisely, however, is not yet clear. Nor is it clear what form future regulatory responses will take. In order to facilitate a dialogue on this important topic, De Nederlandsche Bank (DNB) presents its preliminary view on a set of possible principles in this discussion paper. We believe that the issues and ideas outlined in this paper will benefit from a broader discussion. Particularly relevant questions regard the need for, as well as the scope and form of a potential regulatory response. Furthermore, this paper invites a discussion on what market participants can already do themselves in order to ensure compliance with existing regulations.

DNB, therefore, welcomes comments on this discussion paper. Comments can be sent to ai@dnb.nl. Please note that the deadline for the submission of comments is 18 October 2019. Comments submitted after this deadline, or submitted via other means may not be processed.

All contributions received will be published following the close of the consultation, unless you request otherwise. If you request that your response is treated as confidential, it will not be published on DNB's website, or shared with any third parties.

DNB will use this discussion paper, and the comments received, to engage in a dialogue with the Dutch financial sector over the coming months. We will report on the outcome of this process in the course of 2020.

Summary

Artificial Intelligence (AI), defined as "the theory and development of computer systems able to perform tasks that traditionally have required human intelligence", is becoming progressively more commonplace in finance. This paper provides basic background information on AI, as well as its current and future use in the Dutch financial sector. It concludes with a set of general principles for a responsible application of AI in the financial sector.

AI, as a research area, has been around since 1950. Driven by an exponential increase of computing performance, AI has continuously evolved over time. The period between 1950 and 2010 can be described as 60 years of anticipation, during which steady progress was made, but real breakthroughs – like the 1997 defeat of Garry Kasparov by Deep Blue – were few and infrequent. Since 2010, AI seems to have come of age, and milestones have been achieved almost on a yearly basis. In the last decade, AIs have bested humans in increasingly difficult tasks such as playing complex games – 'Go', poker, and real-time strategy games – and driving cars. Nowadays AI consists of a blooming and diverse landscape of cutting-edge research areas, technologies in various stages of development, and commercial real-world applications. In the near future, AI is expected to become more advanced and increasingly ubiquitous, driven by the continuous increase of computer power, data availability and the advance of the Internet of Things. Various research reports project a 1,000 to 2,000 percent increase in AI-related revenue in the years 2018-2025 and financial services is consistently named as one of the sectors that invests most heavily in AI.

Current applications of AI in the financial sector are manifold and widespread, both in front-end and in back-end business processes. Examples of current AI applications are advanced chatbots, identity verification in client onboarding, transaction data analysis, fraud detection in claims management, pricing in bond trading, AML monitoring, price differentiation in car insurance, automated analysis of legal documents, customer relation management, risk management, portfolio management, trading execution and investment operations. Taking a forward-looking perspective we can expect further AI-driven innovations in various financial domains, ranging from leaner and faster operations ("doing the same thing better") to completely new value propositions ("doing something radically different"). A recent report by the World Economic Forum provides a few dozen examples of potential AI innovations in deposits & lending, insurance, payments, investment management, capital markets and market infrastructure. From an AI-perspective, finance is special in a number of important ways that warrant an adequate regulatory and supervisory response. First of all, the financial sector is commonly held to a higher societal standard than many other industries, and incidents with AI could have serious reputation effects for the financial system. Second, incidents could have a substantial impact on financial stability.

Given the inherent interconnectivity of the financial system, the rise of AI also has a strong international dimension. Furthermore, the specific data environment of the financial sector is a challenge for certain AI applications. Last but not least, the progress of AI on the one hand, and its increasing importance in financial sector business models on the other hand, invites us to rethink our traditional supervisory paradigms.

Based on the observations in this document, and building on the previous work by other regulatory bodies, we have formulated a number of general principles regarding the use of AI in the financial sector. The principles are divided over six key aspects of responsible use of AI, namely (i) soundness, (ii) accountability, (iii) fairness, (iv) ethics, (v) skills, and (vi) transparency (or 'SAFEST').

From a prudential perspective, soundness is the aspect of AI that is of our primary concern. AI applications in the financial sector should be reliable and accurate, behave predictably, and operate within in the boundaries of applicable rules and regulations. This aspect becomes particularly important when financial firms start to apply identical (or relatively similar) AI-driven solutions and systemic risks might arise. Firms should also be accountable for their use of AI, as AI applications may not always function as intended and can result in damages for the firm itself, its customers and/or other relevant stakeholders. Although fairness is primarily a conduct risk issue, it is vital for society's trust in the financial sector that financial firms' AI applications – individually or collectively – do not inadvertently disadvantage certain groups of customers. Financial firms should therefore be able to define what their conception of fairness is and demonstrate how they ensure that their AI applications behave accordingly. As AI applications take on tasks that previously required human intelligence, ethics becomes increasingly important and financial firms should ensure that their customers, as well as other stakeholders, can trust that they are not mistreated or harmed – directly or indirectly – because of the firm's deployment of AI. When it comes to skills, from the work floor to the board room, people should have a sufficient understanding of the strengths and limitations of the AI-enabled systems they work with. Transparency, finally, means that financial firms should be able to explain how and why they use AI in their business processes and (where reasonably appropriate) how these applications function.

1 Introduction

Artificial Intelligence (AI) is becoming more and more commonplace in finance. Financial firms increasingly use AI to enhance their business processes and improve their product and service offerings. The advent of AI has often been described as the Fourth Industrial Revolution.[1] Whether or not AI will live up to this claim, at the very least, it should be considered as an impactful development that requires an adequate response from regulators and supervisors. It should therefore not be a surprise that an increasing number of supervisors and international organisations are publishing their takes on AI and its ramifications for financial sector supervision.

But what do we mean exactly when we talk about AI? Definitions vary widely and have changed over the years. Most financial sector studies by now have adopted the broad definition put forward by the Financial Stability Board (FSB) of AI as "the theory and development of computer systems able to perform tasks that traditionally have required human intelligence".[2] Depending on the nature and objective of future studies, other definitions may be adopted, but for this paper the FSB definition will suffice.

This document should be seen as a first supervisory step into this phenomenon. It aims to provide basic background information on AI (chapter 2), as well as its current and potential future use in the financial sector (chapter 3). Finally, chapter 4 provides a set of general principles for financial firms that apply AI in their business processes.

1 See for instance WEF (2019).

2 FSB (2017, p. 35).


technibel

2 AI: past, present and future

This chapter provides a brief overview of the history of AI and projections regarding its development in the near future, and is meant to provide some background and context for the subsequent chapters. Paragraph 2.1 describes the period of 1950-2009, during which milestones in AI were mainly conceptual in nature, but promises remained largely unfulfilled. Paragraph 2.2 details the coming of age of AI from 2010 to the present day, and paragraph 2.3 provides a forward-looking perspective in which AI will steadily become ubiquitous in our everyday lives.

2.1 Past (1950-2009): 60 years of anticipation

Although AI is a hot topic that seems to have only entered the scene recently, the concept (in its modern form) has been around for at least seven decades. Already in 1950, Alan Turing posited that "machines will eventually compete with men in all purely intellectual fields".[3]

The first electronic digital computers, developed in the late 1930s were little more than gargantuan calculators operating though vacuum tubes. With the replacement of vacuum tubes by transistors in 1947 and the subsequent introduction of the integrated circuit ('microchip') in 1963, computers soon developed into smaller and more powerful machines that started to decorate our desktops in the 1980s. In Figure 1, we can clearly see the exponential growth of computing power since 1945 from the world's first supercomputer 'ENIAC' (capable of 500 floating point operations per second – or 'flops' – see box 1) to today's fastest supercomputer 'Summit' (capable of a staggering 200 petaflops).[4,5] For context, the graph also shows the development of computing power of consumer video cards (graphics processing units) and high-end smartphones.[6] Through the second half of the 20th century, this exponential growth of computing power gave rise to evermore advanced and capable AI applications.

3 Turing (1950, p. 460).
4 Chodosh (2017).
5 See TOP500.org (2019). This record is expected to soon be shattered as AMD and Cray are developing a new supercomputer with 1.5 exaflops (or: 1,500 petaflops) of processing power for the US Department of Energy, to be operational in 2021 (Vincent, 2019).
6 It is often mentioned that today's iPhone has more computing power than NASA's computers that were used to put a man on the moon in 1969. To put this in perspective, as can be seen in figure 1 today's fastest high-end smartphone has about 100,000 times the computing performance as the fastest supercomputer available in 1969.

Box 1 FLOPS

Floating point operations per second – or 'FLOPS' – is the most common metric used to describe computing power when it comes to AI. A floating point operation is any basic mathematical operation that involves decimal numbers (such as '0.007'; as opposed to integers such as '3'). As performance increases exponentially, typical shorthand notations used are kiloflops or KFLOPS (10^3 flops), megaflops (10^6), gigaflops (10^9), teraflops (10^{12}), petaflops (10^{15}), exaflops (10^{18}), and zettaflops (10^{21}).

In the 1950s the first (very small scale) artificial neural networks were developed to emulate the behaviour of a few dozen neurons.[7] The first checkers playing program was also developed during this decade.[8] In 1956 the first academic conference was held on, what organiser John McCarthy coined, 'artificial intelligence'.[9] In 1959 the same McCarthy co-founded the world's first AI lab at MIT.[10]

During the 1960s academic researchers further developed the conceptual framework for AI, amongst others by laying down the foundations of a mathematical theory for AI and introducing the idea of semantic nets.[11,12] The 1960s also saw the light of the first computer-based pattern recognition algorithms and the first attempt at a chatbot called 'ELIZA'.[13,14] In 1968, the first chess playing program was developed that could compete at amateur level, and the world's first autonomous robot 'Shakey' was demonstrated by the Stanford Research Institute.[15,16] 1968 was also the year that AI was popularised (albeit in a negative light) in the persona of HAL 9000, the murderous AI featuring in Arthur C. Clark's '*2001: A Space Odyssey*'.

In the 1970s a lot of work was devoted to further develop the conceptual frameworks for artificial neural networks, natural language processing and visual perception.[17] In 1977, the first backgammon playing program was developed that could compete at a professional level and in 1979 Stanford

7 Yadav, et al. (2015).
8 Samuel (1959).
9 The conference is colloquially called the 'Dartmouth Workshop' (Sarangi & Sharma, 2019).
10 CSAIL (n.d.).
11 Solomonoff (1960).
12 Quillian (1968).
13 Uhr & Vossler (1961).
14 Weizenbaum (1966).
15 Greenblatt, Eastlake, III & Crocker (1967).
16 Stanford Research Institute (n.d.).
17 Driven by, among others, Linnainmaa's work on algorithmic probability (1970), Winograd's SHRDLU program (1971) and the University of Edinburgh work on 'Freddy the Robot' (Artificial Intelligence Applications Institute, 2005).

Figure 1 Computing performance over time

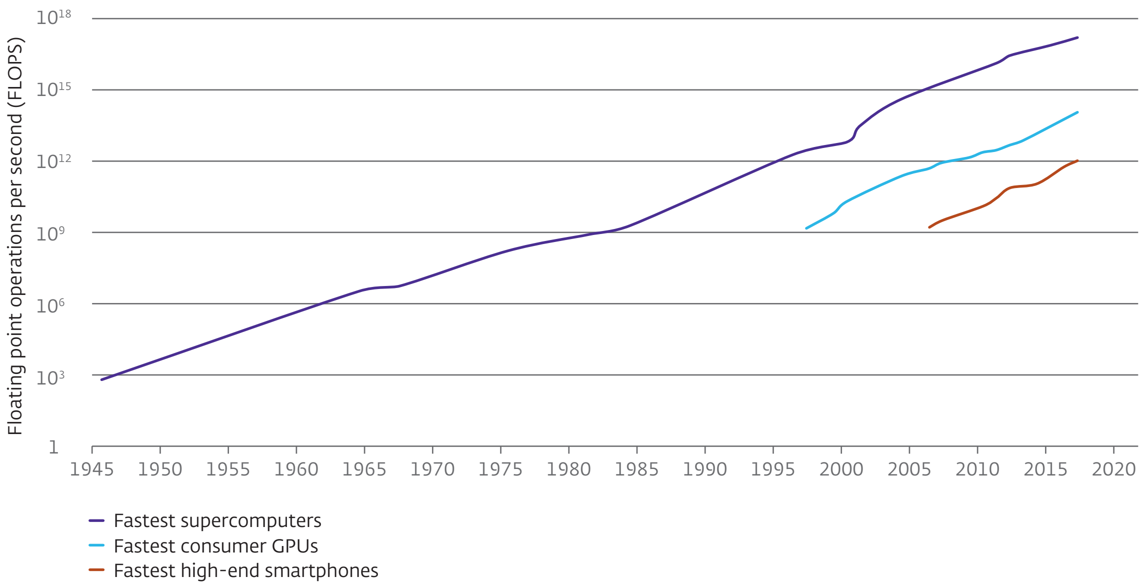

University students built the self-navigating 'Stanford Cart', which can be considered as the ancestor of today's self-driving cars.[18,19] During the 1970s expert systems were also developed, in which knowledge was added to the system by programmers but without any ground-breaking results (leading to the first "AI winter" in which progress slowed down and funding decreased).[20] In the 1980s Explanation Based Learning – the precursor of today's machine learning (see box 2) – was introduced to allow a computer to create a set of rules based on training data.[21] Later that decade, 'NETtalk' was the first successful attempt at using artificial neural networks for learning natural language pronunciation, and in 1989 Tim Berners Lee invented the World Wide Web (which would become instrumental in the development of AI in the 2000s).[22,23]

18 Berliner (1977).
19 Kubota (2019).
20 Examples of such expert systems were Shortliffe's MYCIN program (1977) and so-called LISP-machines (Bawden, et al., 1977).
21 Milton, et al. (1989).
22 Rosenberg & Sejnowski (1986).
23 World Wide Web Foundation (n.d.).

Box 2 Machine Learning

Machine learning comes in various sizes and shapes. The most important categories are listed below.

Supervised Learning
Supervised learning algorithms construct models for predicting outcomes based on past observations. Once the model is trained by being fed example input-output data, it can take inputs only and make output predictions for the purpose of classification or regression. Supervised ML is far more commonly used than unsupervised ML. Using transaction history and patterns to detect fraud, or using loan data and history for credit scoring are examples of supervised learning.

Unsupervised Learning
In unsupervised learning, the model finds patterns based only on input data. Most unsupervised learning implements cluster analysis techniques that split data items into groups, based on the similarity of their features. A common example of unsupervised learning is anomaly detection (eg. transaction monitoring for AML). Another common example is market research, in which unsupervised learning is used to partition the general population of consumers into market segments based on data such as region, spending patterns, or social media activity.

Reinforcement learning
Reinforcement learning refers to goal-oriented algorithms, which learn how to attain a complex objective (goal) or maximise along a particular dimension over many steps; for example, maximise the points won in a game over many moves. Such an AI can start from a blank slate, and under the right conditions, like with DeepMind's 'AlphaGo' (see §2.2), can achieve superhuman performance.

Deep Learning
When there are many hidden layers in an artificial neural network, mimicking the working of a human brain, it is considered a deep learning model. With more layers comes the ability to solve more complex problems, albeit at more computational expense. Deep learning is able to achieve extremely high levels of accuracy, but is also at the root of the "black box" problem in which the human who constructed the model can no longer really understand what the model is doing or how it is making its decisions.

Following the continuing disenchantment with knowledge-based systems, resulting in another "AI winter" in the late 1980s, the focus shifted to data-driven machine learning and reinforcement learning in the 1990s. From there, the first companies started to integrate machine learning in their business processes (like Mellon Bank, which used machine learning to detect credit card fraud).[24]

The 1990s was also the decade in which AIs started to outperform humans in certain areas. In 1994, the reigning checkers world champion Marion Tinsley resigned in a match against the computer program 'Chinook'.[25] Only three years later, in 1997, the world witnessed the defeat of chess grandmaster Garry Kasparov by IBM's 'Deep Blue'.[26] The 1990s also welcomed Google and Amazon on the digital stage,

24 OECD (1997).
25 Madrigal (2017).
26 Weber (1997).

which would become two of the most prominent drivers of today's AI developments.

Driven by increasingly powerful computer systems, most notably in the form of graphics processing units (GPU), AI finally started to get real traction in the 2000s. This progress coincided with the development of deep learning as a subfield of machine learning.[27] The introduction of Sony's PlayStation 3 (PS3) in 2006 also provided a boost for AI research, as the PS3 design and programming allowed researchers to link multiple of these gaming systems together, enabling them to build their own low-cost supercomputers.[28] A year later, Apple introduced the iPhone, setting the stage for mobile internet. Finally, after 60 years of anticipation, the promise of AI would start to come to fruition in the 2010s.

2.2 Present (2010-2019): The coming of age of AI

Since 2010, milestones in the world of AI have been achieved almost on a yearly basis. Besides the ever increasing computer power available to AI researchers, the rise of the World Wide Web in the 1990s and 2000s has created massive amounts of data to feed data-hungry AI systems. As with computing performance, the annual data production continues to grow at an exponential rate.

As can be seen in Figure 2, in 2019 we will create more than 40 zettabytes of data.[29,30] To put this in perspective: if we would write all this data to CD-ROMs, and stack them up neatly in a pile, that pile would be 80,000,000 kilometres high. That is about 1.5 times the distance from Earth to Mars.

The first important milestone for this decade came in 2010 when Microsoft launched Kinect for Xbox 360: an array of sensors that could track human body movement accurately in three dimensions and in real-time.[31] A year later in 2011, IBM's 'Watson' defeated two top contestants in the popular TV show 'Jeopardy!'.[32] That same year, Apple launched 'Siri', the first widely available conversational AI. This was later followed by 'Google Now' in 2012 and Amazon's 'Alexa' and Microsoft's 'Cortana' in 2014. In 2012, Google built a neural network of 16,000 computer processors that successfully taught itself to identify cats in YouTube videos using deep learning.[33] That same year, 'AlexNet' – a convolutional neural network – made headlines at the ImageNet Large Scale

27 Driven by, among others, the work of Geoffrey Hinton (deeplearning.ai, 2018).
28 Moss (2018).
29 A zettabyte is 10^{21} bytes.
30 Reinsel, Gantz, & Rydning (2018).
31 Not only was the Kinect's software developed utilizing machine learning (Microsoft, 2011), the device also facilitated AI research in human motion recognition and lower-level computer vision tasks such as segmentation and object detection (Lun & Zhao, 2015).
32 Markoff (2011).
33 Le, et al. (2012).

IAP Engine Inside
WARNING
IAP®

Figure 2 Annually created data world-wide

Zettabytes (10^{21} bytes)

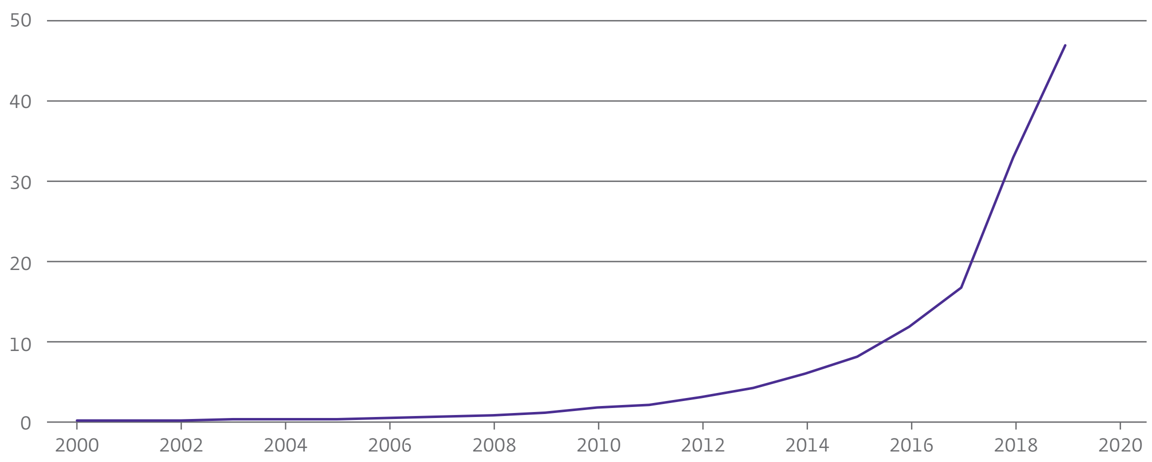

Visual Recognition Challenge, by defeating the competition by a large margin.[34,35]

In 2016, another leap in AI research was achieved when Google DeepMind's 'AlphaGo' spectacularly defeated world champion Lee Sedol in a 5-round game of 'Go'.[36] Go is known as a game that is significantly more complex and difficult to master than chess. A year later in 2017, DeepMind's 'AlphaGo Zero' unceremoniously defeated its predecessor 'AlphaGo' after only 24 hours of learning the game by playing against itself; the score was 100-0.[37] This accomplishment was facilitated by the use of Google's Tensor Processing Units, which are specifically developed for AI applications. Subsequently, its successor 'AlphaZero' successfully applied transfer learning to excel in two other board games: it became the world's best chess player in only 4 hours and the world's best Shogi player in a mere 2 hours.[38] Also in 2017, the poker-playing AI called 'Libratus' managed to defeat top professional players in heads-up games of Texas Hold'em poker.[39] Liberatus became the first AI to achieve a superhuman level of proficiency in a game of imperfect information (with games like chess and 'Go', the information available to the players is

34 A convolutional neural network is a class of deep neural networks.
35 Krizhevsky, Sutskever, & Hinton (2012).
36 Mundy (2016).
37 Silver, et al. (2017).
38 Silver, et al. (2018).
39 Brown N. & Sandholm (2018).

identical or 'perfect', but with a game like poker a player does not have all the information as he does not know the other player's cards). In 2019, Google's DeepMind achieved another milestone when its 'AlphaStar' defeated a pro player in 'StarCraft II'. [40] This computer game has long been considered one of the most challenging real-time strategy games (as opposed to computer games where players take turns to plan and execute their actions). StarCraft II was considered an exceptionally difficult game to master for an AI: it is an imperfect information game, that requires the balancing of short and long-term goals, managing both the macro (economy) and micro (individual unit control) level, and the need to continuously adapt to unexpected situations; all in real-time.

In May 2019, Tesla presented its 'Full Self-Driving' package for autonomous driving.[41] In 2017, Tesla already enabled the 'Enhanced Auto Pilot' mode via a firmware update, making the Tesla the first commercially available (semi) self-driving car.[42] Google's Waymo, Uber, and traditional car manufacturers like Mercedes-Benz and General Motors are also working hard to bring their versions of a self-driving car to market.[43]

The last decade also saw the introduction of relatively easy to use AI-specific software packages and low-cost cloud-based solutions that enabled researchers, corporations and even SMEs to develop their own AI-driven applications.[44] This proliferation of AI capabilities, however, also brought to light certain risks associated with AI. As Bill Gates once said: "it is fine to celebrate success but it is more important to heed the lessons of failure". Box 3 below therefore expounds some of AI's cautionary tales.

Today the world of AI consists of a blooming and diverse landscape of cutting-edge research areas, technologies in various stages of development, and commercial real-world applications. Figure 3 below presents an overview of various AI technologies and research areas. An in-depth discussion of all technologies and research areas mentioned falls outside the scope of this paper, but the overview does show that the field of AI and its potential future applications in the financial sector entail a lot more than 'just' predictive analytics and machine learning. The next paragraph will provide a few reflections as to what we may expect regarding the development of AI in the coming decade.

40 DeepMind (2019).
41 Rapier & Bastone (2019)
42 Brown, M. (2017).
43 Welch & Behrmann (2018).
44 Most notably the AI-related services available on Microsoft Azure Machine Learning Studio, Google Cloud Machine Learning, and Amazon Web Services SageMaker.

Box 3 When AI gets it wrong: errors, accidents and disasters

American Economist Leo Cherne is attributed with saying that "The computer is incredibly fast, accurate, and stupid". Despite the reference to 'intelligence' in AI, this holds true for today's AI as well. Current AI applications are exceptionally strong at performing narrow, specific tasks. However, they lack the general intelligence to reflect on their performance, and typically fail to catch errors that are plain obvious to humans. Some of these errors are merely funny, others can be painful, damaging, or even catastrophic.

One reason AI fails, is because it has not been properly trained for all eventualities. In November 2018, an AI system in the southern port city of Ningbo, China, falsely classified a photo of Chinese billionaire Mingzhu Dong on an ad on a passing bus as a jaywalker. This embarrassing mistake went viral on Chinese social media and Ningbo police were quick to apologise.[45] In June 2016, a Tesla driver was killed while the car was in auto pilot mode, when it crashed into a truck. After careful examination of the logs, Tesla stated that "neither Autopilot nor the driver noticed the white side of the tractor trailer against a brightly lit sky, so the brake was not applied".[46]

Another reason for AI failure is poor training data. A 2016 article from the University of Washington demonstrates how an artificial neural network that was trained to distinguish between wolves and huskies failed as a result of poor training data. Because all of the training images of wolves had snow in the background, and those of the huskies did not, the AI did not distinguish between features of the animals themselves but simply used the background as the discriminating feature.[47] Such accidents, however, are not confined to academic experiments. In 2015, Google got embarrassingly called out when its Photos app labelled black people as "gorillas".[48] Just a month before this incident, Yahoo's Flickr app misclassified pictures of a concentration camp as "jungle gym" and "sport".[49]

Sometimes AIs get it wrong, simply because they adopt societal biases. In 2018, Amazon had to scrap its AI-enabled recruiting software because it was biased against women. This bias was a direct reflection of the male dominance in the tech industry, and even after several attempts to exclude gender as a feature, the AI kept finding proxies for gender to discriminate on.[50] Another well-known example is how popular online translation engines struggle to be gender-neutral. For example, if you would translate "she is a doctor" from English to Turkish, the result is "o bir doctor". Translate it back, however, and because of gender bias in English literature translation engines typically return "he is a doctor".[51]

AIs are also prone to be corrupted by malicious intent. In 2016, The Verge headlined "Twitter taught Microsoft's AI chatbot to be a racist a**hole in less than a day". Microsoft's Twitter bot 'Tay' was caught twittering very explicit racist and hateful messages within 24 hours after it went live.[52]

People, of course, are fallible and prone to failure too. However, when an AI fails, it fails diligently, at scale, and often in unexpected ways. This in particular is the reason, at least for AI applications in finance, why it is so important to get things right from the start.

45 Dodds (2018).
46 Boudette (2017).
47 Ribeiro, Singh & Guestrin (2016).
48 Grush (2015).
49 Hern (2015).
50 Dastin (2018).
51 Sonnad (2017).
52 James (2016).

Figure 3 Overview of AI technologies and research areas

Artificial general intelligence
Transfer learning

Automation
Automated planning and scheduling
Automated reasoning
Automatic target recognition

Biologically inspired computing
Evolutionary computation
Neuromorphic hardware

Computer audition
Speech recognition
Speaker recognition

Computer vision
Image processing
Intelligent word recognition
Object recognition
Optical mark recognition
Optical character recognition
Facial recognition systems

Expert system
Decision support system

Game artificial intelligence
Computer game bot
Video game AI
General game playing
Game theory

Intelligent agent
Agent architecture

Knowledge management
Concept mining
E-mail spam filtering
Information extraction
Information extraction
Named-entity extraction
Knowledge representation
Semantic web

Machine learning
Constrained Conditional Models
Supervised learning
Unsupervised learning
Reinforcement learning
Deep learning
Neural modeling fields

Natural language processing
Chatbots
Language identification
Natural language user interface
Natural language understanding
Machine translation
Question answering
Semantic translation

Pattern recognition
Optical character recognition
Handwriting recognition
Speech recognition
Face recognition

Robotics
Behavior-based robotics
Cloud robotics
Cybernetics
Developmental robotics
Epigenetic robotics
Evolutionary robotics

Speech generating device
Natural language generation

Various
Affective computing
Argument technology
Artificial creativity
Artificial life
AI diagnosis
Hybrid intelligent system
Intelligent control
Nonlinear control
Strategic planning
Vehicle infrastructure integration
Virtual Intelligence

Virtual reality
Augmented reality

2.3 Future (2020-beyond): Ubiquitous AI

As with the past 70 years, the main driver for future developments will be the further increase of computing power. Figure 4 shows the projected development of computing performance from 2020-2035. Compared to the fastest supercomputers, GPUs in home PCs have a lag of about 12 years in terms of computing performance. Subsequently, mobile phones lag about 9 years behind home PCs. In other words, if the trend of the last few decades continues, the performance of today's fastest supercomputer will sit on your desk at home by 2030 and will fit comfortably in your jeans' pocket by the year 2039 (see box 4 for some further discussion regarding the future development of computing performance).

What is perhaps even more thought-provoking is that, in terms of raw computational power, we are likely to achieve (and then quickly surpass) the computational power equivalent of the human brain in the coming years. Although our brains do not perform floating point operations as computers do, researchers estimate that the human brain has a performance equivalence in the exa-zettaflops range.[53] Now, achieving human brain equivalent computing power, in and of itself, does not mean that we will suddenly have created a sentient entity with human-like general intelligence. The human brain is a very complex, and highly efficient piece

Figure 4 Projected computer performance

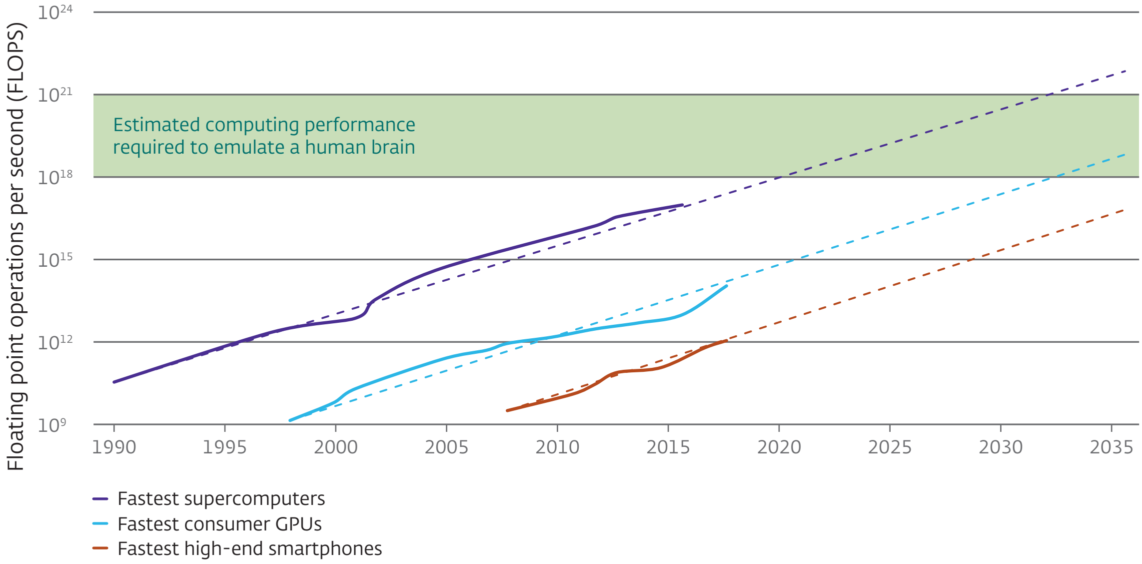

53 Bouchard, et al. (2018), Sanberg & Bostrom (2008).

of biological hardware, and neuroscience still has a long way to go in fully understanding its workings. It does raise the question what kind of feats AI applications may accomplish in the next decade, given that current AIs are already able to beat humans in increasingly complex cognitive tasks.

Another, perhaps more relevant, way of looking at all this is the implication that currently available feats of AI will come in increasingly smaller packages. This means that AI applications that presently need a lot of expensive hardware or cloud-services to operate will become available to desktop environments and, eventually, in tiny devices, as time progresses. The 'Internet of Things' is a direct consequence of this trend, which will gradually bring AI into all aspects of our daily lives, quietly operating in the background and creating ever more data in the process.[54] Mark Weiser, Chief Technologist at Xerox PARC in the 1990s, said that "the most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it".[55] This quote aptly captures what we may expect from AI in the coming years from a consumer-perspective. Following the self-driving car, we will likely see the advent of other automated services, like 'self-driving finance' in which AI will first augment, and eventually autonomously manage, our finances.[56]

From an industry-perspective, former Google and Baidu scientist Andrew Ng describes AI as the new electricity: "Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don't think AI will transform in the next several years".[57] Ng regards financial technology as one of the sectors best suited for an AI-led transformation, and states that the only factors that could slow down this development are the scarcity of talent, and the challenge for companies to effectively harness their data.

Various research reports project a 1,000 to 2,000 percent increase in AI-related revenue in the years 2018-2025.[58] Financial services is consistently named as one of the sectors that invests most heavily in AI, and in AI-related big data and business analytics, banking is even named as the largest revenue-generating sector.[59,60] The next chapter provides a more in-depth overview of current and potential future AI applications in the financial sector.

54 Brunkhard (2018).
55 Weiser (1991, p. 94).
56 WEF (2018).
57 Lynch (2017).
58 Tractica (2019), Orbis Research (2017), Grand View Research (2017).
59 Statista (2016), Ibid. (2015).
60 Ibid. (2019).

Box 4 The end of Moore's Law?

the promise of exponential growth of computing performance was more or less formalised with "Moore's Law" (named after Intel's co-founder Gordon Moore) in 1965. Moore projected that the number of transistors in a dense integrated circuit would continue to double about every two years. As we have seen in paragraph 2.1, the exponential increase of computing performance has kept a steady pace during the last 70 years.

Moore's Law, however, is not a law of nature, but rather a prediction about future technological developments that has been a guiding principle for major chip-developers like Intel, NVIDIA and AMD. This means that historical growth-rates do not guarantee that these rates will be sustainable in the future. In the last few years chip-developers have struggled to make the step from a 10 to a 7 nanometre lithography process in the production of computer chips, and in 2015, Intel admitted that it would no longer be able to keep up with Moore's Law. Instead of doubling transistors every 24 months, Intel said, this process would advance every 30 months from then on.[61] This prompted many commenters to proclaim the imminent end of Moore's Law. Although the challenge of creating ever-smaller transistors undeniably becomes harder, there are good reasons to expect that computing performance will continue to grow exponentially in decades to come.

First of all, computer performance has increased at a consistently higher pace than the growth rate that characterises Moore's Law, doubling every 18 months instead of every 24 months. This is because Moore's Law is about the *number of transistors* only, and does not take into account innovations that increase the *performance* of the chips (such as higher clock speeds, multi-processors and multi-threading). In 2018, NVIDIA CEO Jensen Huang underlined that the number of transistors is of diminishing importance and that, if anything, performance increase is speeding up.[62] Furthermore, technologies and methods are being developed to circumvent the challenges of even lower nanoscale lithography processes on silicon, for example by moving from a single layer (2D architecture) to multilayer (3D architecture) computer chips, the adoption of carbon nanotubes technology to replace silicon, quantum computing, and bio-computing.[63, 64, 65, 66] Finally, even if performance increases would slow down at some point, this would not change the exponential nature of this increase, but would only mean that new milestones, such as reaching zettaflops-performance, will be achieved at a somewhat slower (but still exponential) rate.

61 Oreskovic (2016).

62 Perry (2018).

63 Giles (2018).

64 Malecek (2016).

65 Bachor & Simmons (2017).

66 Nicolau, et al. (2016).

Last 2268.00
Prev Close 2259.00
10 Oct 2011
Graph 60
GX1 Index
Options
07 Oct 2011
10 Oct 2011
Graph 64
PAA Comdty
Last 611.00
Prev Close 585.85
Prev Close 82.98
Graph 63
MF1 Index
PLA Comdty
Rabobank:
De Telegraaf
Het Fin Dag
Nedap: RF
Philips Ele
DNB: Auto
RTLZ:
DePers:
Dutch St
Ahold
DELL

3 AI in finance

This chapter will provide some insight into how AI is currently already being used in finance, and what potential applications we might expect in the near future. Paragraph 3.1 presents an overview of current use cases in the financial sector, paragraph 3.2 describes potential future applications, and paragraph 3.3 analyses specific AI-related issues from a financial sector perspective.

3.1 Current applications

AI adoption in the financial sector already is quite widespread, even though we might not always be aware of its presence. In their report "Artificial Intelligence and machine learning in financial services" the FSB distinguishes four different areas for AI application: customer-focused applications, operations-focused uses, trading and portfolio management, and regulatory compliance and supervision.[67] The French *Autorité de Contrôle Prudentiel et de Résolution* (ACPR) has observed that the implementation of AI seems more advanced in the banking sector than in the insurance sector.[68] As stated in their report, most applications currently in production deal with operational processes (eg. using NLP to automate the processing of written documents, fraud detection), and so far AI is not yet commonly used for activities that are material to customers (such as credit scoring, advice, and underwriting). Areas of applications provided by other reports are chatbots, claim management, personalised pricing, financial and operational risk management, compliance, anti-money laundering, cyber security, portfolio management, and internal risk modelling.[69]

A number of real-world cases, noting the use of AI in both front-end and back-end business processes of financial institutions, are detailed below. Note that these examples are not exhaustive, their inclusion neither implies that DNB endorses these applications, nor have we verified that these examples adhere to the principles set out in this paper. These cases merely intend to illustrate the manifold contemporary applications of AI in finance.

Customer interaction

An increasing number of financial service providers employ AI-enabled chatbots that use natural language processing techniques to provide a 24/7 first line of customer interaction. Intelligent chatbots can save costs dramatically as they reduce the need for customer service call centres. The Dutch insurance company ASR's chatbot-based website is a typical example of an intelligent chatbot.[70] Increasingly, companies opt to integrate their chatbots into existing platforms such as Facebook Messenger or WhatsApp, as well as

67 FSB (2017).
68 ACPR (2018).
69 FSB (2017), BaFin (2018), Monetary Authority of Singapore (2018), ACPR (2018), CSSF (2018).
70 ASR (n.d.).

speech-enabled platforms such as Google Assistant, Apple Siri and Amazon Alexa.

Client oboarding

Nowadays, when we want to open a new bank account, we are likely to encounter AI during the onboarding process. Rabobank, for instance, uses UK-based iProov's deep-learning driven 'Face Verifier' to match a selfie with the face image from the customer's identity document, including liveness check. Rabobank uses this service in conjunction with 'ReadID' to verify the ID document's authenticity, which also uses AI techniques such as optical character recognition.[71]

Account information services

AI is also commonly used in account information services, such as ABN AMRO's 'Grip' app.[72] Grip uses machine learning to automatically classify incoming and outgoing transactions and assigns them to categories that are relevant to the app's user. The app also provides insight into possible saving opportunities as well as more general budgeting services. With the recent implementation of PSD2, more such apps by third-party service providers are expected to be introduced.

Fraud detection

Dutch fintech company FRISS is an AI-enabled service provider that is used by various Dutch insurance companies for fraud detection in claims management.[73] Their system uses expert knowledge, proprietary data, and external data sources to provide real time risk scores. FRISS also offers AI-driven risk assessment at underwriting and CDD solutions for onboarding new clients.

Investor services

ING's 'Katana' tool helps bond traders with their pricing decisions. Katana employs predictive analytics, based on historic and real-time data, to help traders decide what price to quote when buying and selling bonds for their clients.[74] This tool is a typical example of how AI can be used for augmentation (rather than for automation) as it analyses and visualises relevant information for the user to decide on.

Transaction monitoring

BusinessForensics' HQ Craft is an example of an AI-driven solution for real-time AML transaction monitoring. This big data processing platform uses data analysis of network transactions, forensic analysis, fraud management and machine learning. This Netherlands-based company provides its services to various Dutch banks and insurance companies.[75]

71 Wegdam (2018).
72 ABN AMRO (n.d.).
73 FRISS (n.d.).
74 Ibid. (2017).
75 BusinessForensics (n.d.).

Car insurance
Fairzekering is a Dutch AI-enabled car insurance product. The insurance company analyses their customers' driving behaviour based on the data collected by a smart module 'Chipin' that is installed in their customers' cars. Depending on your risk score, you get a 0-35% discount on your monthly rate. Customers also have access to a dashboard that provides them with additional insights in their driving behaviour.[76]

Document analysis
JPMorgan Chase developed an AI-driven Contract Intelligence platform called 'COiN' to analyse legal documents and extract important data points and clauses. According to JPMorgan, the system was used to review 12,000 annual commercial credit agreements, which would normally take 360,000 hours to review manually.[77] A number of Dutch banks and insurance companies have indicated that they are also implementing automated solutions to analyse documents.

Customer relationship management
Salesforce is an AI-enabled CRM platform that offers companies a wide variety of tools to support their client interaction. One of their products, 'Einstein', uses AI to provide lead scoring, activity capture, opportunity insights, and account insights services. Various Dutch financial firms are among Salesforce's customers.[78]

Portfolio management and risk analytics
BlackRock's 'Aladdin' is an example of a financial platform that increasingly applies AI.[79] Aladdin is a widely used operating system that combines AI with sophisticated risk analytics, comprehensive portfolio management, trading execution and investment operations tools on a single platform. Today, Aladdin is used to manage an estimated 18 trillion US dollars in assets globally.[80] BlackRock's platform is used by numerous Dutch financial institutions, including major banks and pension funds.

As the examples above show, the use of AI is already widespread in the financial sector, and we can expect that AI will be used increasingly to augment or automate existing business processes as well as to enable new financial products and services.

3.2 Potential applications in the near future

Various reports provide a glimpse of what AI has to offer the financial sector.[81] AI applications can for instance be used to structure relatively complex middle and back office settlement processes

76 Fairzekering (n.d.).
77 JPMorgan Chase & Co (2017).
78 Salesforce (n.d.).
79 Kochansky (2018).
80 Massa (2018).
81 FSB (2017), BaFin (2018), Monetary Authority of Singapore (2018), ACPR (2018), CSSF (2018).

more efficiently or even automate them. AI can furthermore improve the customer experience and compliance processes (especially in the prevention of money laundering). It can increase the precision of risk assessment models and generally bring efficiency and effectiveness gains to financial firms' core processes. AI can be used to monetise customer data, resulting in new income sources for financial firms. It can simplify and speed up interactions with consumers and result in faster processes in trading, sales and product development. Finally, the application of AI may lead to highly individualised products and services. Last year, the World Economic Forum (WEF) published a comprehensive report titled "The New Physics of Financial Services", which provides an insightful forward-looking perspective.[82] The report gives examples in various financial domains, ranging from leaner and faster operations ("doing the same thing better"), to completely new value propositions ("doing something radically different"). Although it is hard to predict how AI will exactly shape the financial sector landscape in years to come, the examples provided by the WEF report illustrate the wide variety of potential innovations nicely.

Deposits and lending

According to the WEF report, AI can be expected to be increasingly used in deposits and lending to (i) provide just-in time lending, (ii) miniaturise unsecured lending to be use-specific, (iii) provide personalised practicable advice in real time, (iv) offer tailored 'always-on' experiences across channels, (v) predict defaults with greater accuracy, (vi) automate and augment business credit decision-making, and (vii) improve client advisory by integrating into data streams for opportunity discovery.

Insurance

Regarding insurance the WEF foresees AI applications to (i) improve underwriting, pricing efficiency and accuracy, (ii) increase capital efficiency through better risk modelling and real-time risk monitoring, (iii) triage and grade claims to increase adjudicator efficiency, (iv) increase the capabilities of sales agents and advisors, (v) improve scale efficiencies to enter new markets, (vi) develop modern, mobile-first insurance offerings, (vii) reduce fraud using new tools and data, (viii) process claims instantly, (ix) advise clients on prevention strategies to lower their risk profiles, (x) provide predictive analytics to clients that help them better understand their risks, (xi) introduce new pricing and payment models, (xii) develop modularised policies, and (xiii) use proxy data to insure new risk categories.

Payments

The WEF report projects that the payments industry will apply AI to (i) further automate compliance and reporting, (ii) drive loyalty by offering bespoke incentives and rewards, (iii) compete to become a provider of invisible payments infrastructure, (iv) increase detection precision to eliminate false positives, (v) deploy real-time surveillance

82 WEF (2018).

capabilities, (vi) create an advisory capability for macroeconomic trends, (vii) offer 'prediction as a service' to merchants, and (viii) act as the ultimate personal shopper for customers.

Investment management
AI can be deployed in investment services in order to (i) offer seamless account setup and customer acquisition, (ii) control back-office compliance costs, (iii) digitise customer servicing, (iv) equip advisers with highly personalised insights, (v) holistically understand investor preferences in real-time, (vi) offer outcome-based portfolio modelling, (vii) use existing but unused platforms for distribution, (viii) mimic advanced strategies for low costs, (ix) establish passive products that track new datasets, (x) analyse vast quantities of data at scale, (xi) share more detailed economic insights, (xii) continuously source new and exclusive datasets, (xiii) develop unique strategies and new investment products, (xiv) find new and unique correlations between datasets, and (xv) enable users to effectively manage their investments.

Capital markets
Applications in capital markets are foreseen by the WEF to (i) automate pre-trade analysis, (ii) automate investment monitoring and reporting, (iii) use predictive models to improve deal identification, pairing and sales activities, (iv) achieve better investments performance by using new data in opaque markets, (v) develop real-time pre- and post-trade risk-management solutions, and (vi) use broader and better data to develop predictive risk models that drive capital savings.

Market infrastructure
For the sixth and final sector, market infrastructure, the WEF expects AI to be used to (i) automate alert triage, investigation and reporting, (ii) automate reconciliation and incident reporting, (iii) integrate post-trade workflows to achieve straight-through processing, (iv) improve trade speed and price using dynamic execution methods, (v) deploy new order types to protect investors from risks proactively, (vi) deploy holistic market-surveillance services, (vii) develop macroeconomic indicators using internal data, (viii) offer insights on market structure and risk, and (ix) sell internal 'analytics capabilities as a service'.

It is hard, if not impossible, to predict whether and how these global trends described by the WEF and other reports will materialise in Europe and in the Netherlands. To a large extent, this will also depend on future regulatory responses.[83] The new Payment Service Directive (PSD2), for instance, is expected to enable a number of AI-driven innovations, most notably in deposits, lending and payments. The recently implemented General Data Protection Regulation (GDPR), on the other hand, sets explicit and implicit constraints on potential innovations by regulating the manner in which companies may use personal data in their product and service offerings.

83 The European Commission is currently looking into AI, and the Commission's High Level Expert Group is expected to come up with policy proposals later in 2019.

Within Europe, the Dutch financial market seems to be well-suited to become a front-runner in the adoption of AI. Amsterdam and Eindhoven are positioning themselves as upcoming AI hubs and ING has recently launched a new AI Fintech Lab in collaboration with Delft University of Technology.[84, 85]

3.3 What makes finance special

The challenges provided by an increasing adoption of AI are not exclusive to the financial sector, and important questions regarding ethics, privacy and security are relevant for financial and non-financial firms alike. Still, from an AI-perspective finance differs from other industries in a number of important ways that warrant an adequate regulatory and supervisory response.

One such difference is that the financial sector is commonly held to a higher societal standard than many other industries, as trust in financial institutions is considered essential for an effective financial system. Furthermore, many consumers regard financial institutions like banks, insurance companies, and pension funds to be providers of a utility function. The data that these companies extract from their customers are therefore regarded as a by-product of this function, instead of a commercial asset for the financial sector to monetise. Even data usage that is acceptable from a strictly legalistic perspective can result in material reputational damage if it does not correspond to the expectations of society at large.

Another aspect that makes finance special is the potential impact of AI on financial stability. The financial system is not impervious to shocks. As experienced in the global financial crisis, problems can propagate through the system rapidly. When AI increasingly manages financial transactions, relatively small incidents could have major consequences. Given the increasing importance of tech giants in providing AI-related services and infrastructure, the concept of systemic importance may also need to be extended to include these companies at some point.

Given the inherent interconnectivity of the financial system, the rise of AI has a strong international dimension.[86] An adequate policy response will require close international cooperation and clear minimum standards and guidelines for the sector to adhere to. Regulatory arbitrage in the area of AI could have dangerous consequences and should be prevented where possible.

Another way in which finance can be different from other industries is that for many financial services we have a very specific data environment. Because of cultural and legal differences data are, in many cases, only representative for domestic markets.

84 ING (2019).
85 Municipality of Amsterdam (2018)
86 FSB (2017).

Good examples of this are data pertaining to credit, pensions and mortgages. This may provide a challenge for the development of data-hungry AI systems, especially for relatively small markets as that of the Netherlands. Furthermore, because of continuous changes to the financial regulatory framework, historic data quickly becomes less representative and unusable for training AI-enabled systems.

The progress of AI in beating humans in increasingly difficult games should also give us pause. In its design, the stock market is set up as a perfect information game, in which all participants have access to the same information.[87] Banking and insurance can arguably be regarded as imperfect information games: customers have information regarding their financial position and risk profile that is unavailable to banks and insurance companies.[88] As we have seen in the previous chapter, AIs like AlphaGo and Liberatus are making huge strides in beating humans in both perfect and imperfect information games. Now that computer performance is progressing to the point where it equals the raw performance capabilities of the human brain, we should consider the possibility that at some point – perhaps sooner than we think – AIs could surpass human agents in trading, banking and insurance as well.[89]

Finally, the increasing importance of AI for the financial sector also invites us to rethink our traditional supervisory paradigms. When it comes to AI, strong operational controls could be more important than capital and liquidity buffers. We will increasingly come across business models that are driven primarily by the monetization of data, rather than by product fees and interest margins.

87 This property of the stock market is commonly referred to as the efficient-market hypothesis.
88 This is also referred to as the 'lemon problem' (Akerlof, 1970).
89 Algorithmic trading already accounts for about 70% of total orders in some markets. These systems, however, do not currently outperform human traders in the quality of their trade decisions per se, but primarily in response and execution times.

4 General principles for the use of AI by financial institutions ('SAFEST')

As the previous chapters illustrate, AI is a ubiquitous phenomenon. On the one hand, this implies that many of the prerequisites for a responsible application of AI are already covered by existing regulations. In particular the requirements regarding sound and controlled business operations (see Financial Supervision Act articles 3:10, 3:17 and 3:18 for legal context) cover many aspects.[90] On the other hand it is challenging to get a comprehensive overview of the AI-specific challenges faced by financial firms. Moreover, loopholes in regulation may exist and societal expectations may go beyond what current regulations require. Consequently, various legislators and supervisory authorities, including the European Commission and the European Banking Authority, are considering the development of AI-specific policies and/or regulations.[91] We cannot predict, however, how future regulatory frameworks regarding the application of AI in the financial sector will evolve, or the time it will take for these regulatory responses to take shape.

The principles presented in this chapter aim to provide a framework that helps financial firms to assess how responsible their application of AI is and also provide a way to facilitate our own assessment of the direction and desirability of future regulatory developments. It goes without saying that the formulated principles do not preclude that other, more stringent, regulations and/or regulatory practices may already apply to specific regulated activities; in these cases the most stringent regulations take precedence.

The principles in this chapter should be seen in the context of controlled and sound business operations. Proportionality applies to these principles, and their applicability should be considered in light of the scale, complexity and materiality of an organisation's AI applications. The applicability of these principles is also determined by the role of an AI application in the organisation's

Figure 5 Applicability of principles in light of materiality and purpose of the AI

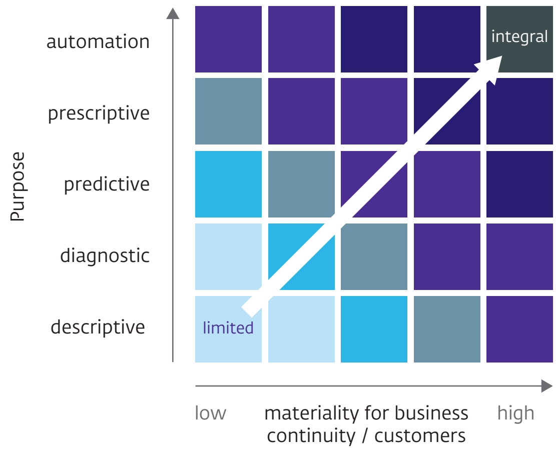

90 Appendix 'A' provides an unofficial translation for the relevant articles.

91 European Commission (n.d.), High-Level Expert Group on Artificial Intelligence (2019), EBA (2019).

decision making process; this means whether the AI application serves a descriptive, diagnostic, predictive, prescriptive, or automation purpose. We can visualise this proportionality in a 'heat map', as can be seen in Figure 5.

The principles are divided over six key aspects of responsible use of AI, namely soundness (§4.1), accountability (§4.2), fairness (§4.3), ethics (§4.4), skills (§4.5), and transparency (§4.6).[92] For each principle suggestions are provided on how the principle can be operationalised in an organisation.

4.1 Soundness

From a prudential perspective, soundness is the aspect of AI that is of DNB's primary concern. AI applications in the financial sector should first and foremost be reliable and accurate, behave predictably, and operate within in the boundaries of applicable rules and regulations, including non-financial regulations, such as the GDPR. This aspect becomes particularly important when financial firms start to apply identical (or relatively similar) AI driven solutions and systemic risks might arise. Financial firms applying AI in their business processes should demonstrate that they have taken all necessary measures to ensure their business continuity.

1) Ensure general compliance with regulatory obligations regarding AI applications.
2) Mitigate financial (and other relevant prudential) risks in the development and use of AI applications.
3) Pay special attention to the mitigation of model risk for material AI applications.
4) Safeguard and improve the quality of data used by AI applications.
5) Be in control of (the correct functioning of) procured and/or outsourced AI applications.

Ad 1: Where possible and appropriate, compliance with applicable regulatory frameworks is implemented in the design of the AI applications itself ("compliance-by-design"). In case of malfunctioning AI systems, fall-back plans are available for core business processes to ensure continuity of operations and regulatory compliance.

Ad 2: Domain experts are actively involved in the development and implementation of AI applications (eg. regarding an AI application to better predict probability of default levels, experts in credit risk and special asset management can be considered relevant domain experts). Where appropriate, boundaries are set to constrain model outcomes, serving either as floors/ceilings or as a mechanism to delegate further decision-making to a domain expert. The choice for and definition of relevant evaluation metrics (such as accuracy, precision,

92 These key aspects are partly inspired by the 'FEAT principles' from the Monetary Authority of Singapore (2018).

recall and specificity) are well substantiated and documented. Models, where appropriate, are periodically retrained, recalibrated and assessed, especially in the event of significant changes in the input data, relevant external factors, and/or in the legal or economic environment. Criteria for significant change as well as other fail criteria are well documented for each AI application, and become more stringent as the model's materiality increases. Procedures are in place for personnel to address issues with the choice of models, used data sets, and/or model outcomes.

Ad 3: Criteria for model choices include considerations other than quantitative evaluation metrics, such as the explainability, simplicity, and reliability of the chosen models. Consistency should be maintained, where reasonable, between models used for similar analyses (eg. between models used for loss estimates and models used for pricing). The impact of incorrect model outcomes for the organisation are periodically assessed. Material AI enabled systems are designed with 'human-in-the-loop' and/or 'human-on-the-loop' review processes. The correct workings of material AI applications is regularly assessed by evaluating model outcomes against conventional (non-AI enabled) models.

Ad 4: Minimal requirements regarding data quality are defined. Efforts are made on a continuous basis to ensure that data are correct, complete, and representative. Special attention is paid to missing or incorrect data-points, potential sources of bias in data, features and inference results (such as selection and survival bias). Procedures and safeguards are in place to maintain and improve data integrity and security during the process of data collection, data preparation, and data management. Issues with data integrity and bias, both in development and production, are evaluated and documented in a structural manner for future reference. Original data sets used to (re)train and (re)calibrate models are systematically archived.

Ad 5: The organisation's AI policy (and associated standards and requirements) is applied consistently to both applications developed in-house as well as applications developed by third parties. Financial firms need to understand and be in control of the risks associated with their use of AI applications, whether these applications are procured or outsourced, developed in a partnership, or have any other material involvement from third parties.

4.2 Accountability

The workings of AI applications can be complex and difficult to understand. Furthermore, AI applications may not always function as intended and can result in damages for the firm itself, its customers and/or other relevant stakeholders. Especially when AI applications become more material, financial firms should demonstrate unequivocally that they understand their responsibility for AI applications and that they have operationalised accountability for these applications throughout their organisation. Model complexity or third party reliance should never be used as arguments for limiting the organisation's accountability.

6) Assign final accountability for AI applications and the management of associated risks clearly at the board of directors level.
7) Integrate accountability in the organisation's risk management framework.
8) Operationalise accountability with regard to external stakeholders.

Ad 6: The board of directors understands their responsibilities regarding the adoption and use of AI, and takes genuine responsibility for the risks associated with their AI applications. Operational accountability is explicitly assigned at all relevant levels of the organisation, and final accountability for AI applications and their outcomes (both for the organisation and their customers) is assigned to one (or more) board member(s). The board of directors accountability explicitly extends to externally developed and/or sourced AI applications.

Ad 7: The adoption and use of AI is integrated in the organisation's risk management framework. Clear roles and responsibilities are assigned throughout the organisation to ensure the responsible use, management and auditability of AI applications.

Ad 8: On request (eg. by a customer or an auditor) and where appropriate, specific AI driven decisions and/or model outcomes are reviewed by a domain expert in order to verify (or adjust) and explain these decisions/outcomes. Verified and relevant supplementary data provided by customers are taken into account when performing a review of AI driven decisions.

4.3 Fairness

Although fairness is primarily a conduct risk issue, it is vital for society's trust in the financial sector that financial firms' AI applications – individually or collectively – do not inadvertently disadvantage certain groups of customers. Whether we are talking about an AI that monitors transactions to detect money laundering, or one that facilitates loan decisions, financial firms should be able to define their concept of fairness and demonstrate how they ensure that their AI applications behave accordingly.

9) Define and operationalise the concept of fairness in relation to your AI applications.
10) Review (the outcomes of) AI applications for unintentional bias.

Ad 9: Fairness is taken into account in the design and training of the AI application, both in the selection of input parameters (especially when these contain personal attributes) and potential sources of unintentional bias ("process fairness"; preventing disparate treatment). Fairness is qualitatively defined in terms of group fairness and/or individual fairness. Based on the definition, it is subsequently operationalised in evaluation metrics, for example in terms of false positive and/or false negative rates ("outcome fairness"; preventing disparate impact). Trade-offs between process fairness, outcome fairness and accuracy (soundness) are substantiated, well documented and available for review.

Ad 10: Material AI enabled systems (in terms of potential impact for customers) are designed with 'human-in-the-loop' and/or 'human-on-the-loop' review processes to detect and minimise unintentional bias. Procedures are also in place for 'after the fact' reviews, on instigation of the customer and/or other stakeholders.

4.4 Ethics

As AI applications take on tasks that previously required human intelligence, financial firms should ensure that the outcomes of these systems do not violate the firm's ethical standards. This moral obligation goes above and beyond compliance with applicable legal requirements. Financial firms should ensure that their customers, as well as other stakeholders, can trust that they are not mistreated or harmed – directly or indirectly – because of the firm's deployment of AI, even if such applications operate within the boundaries of applicable laws and regulations.

11) Specify objectives, standards, and requirements in an ethical code, to guide the adoption and application of AI.
12) Align the (outcome of) AI applications with your organisation's legal obligations, values and principles.

Ad 11: The organisation's policy includes criteria that inform consistent decision making on whether specific processes and functions are (or are not) considered suitable for deployment of AI. The policy includes criteria that inform decision making on the use (or disuse) of specific models, methods, and data in AI applications. Procedures are in place that channel ethical and other material concerns that may arise from the development, implementation, and use of AI to an ethics commission or other suitable organisational body.

Ad 12: Criteria for the fitness of model outcomes are determined and documented a priori, and are informed both by legal requirements and by the institution's values and principles. In the design and approval of customer-oriented AI applications, due consideration is given to the customer's interest. AI applications do not (inadvertently) exploit consumer's behavioural patterns or psychological biases in a way that is harmful to their (financial) well-being.

4.5 Skills

As most AI applications augment human tasks, financial firms' employees and management will increasingly come to rely on AI applications to support them in their work. Like with many other tools, wrong use can result in accidents and it is the organisation's responsibility to ensure that their (senior) management, risk management and compliance functions have an adequate level of expertise. From the work floor to the board room, a sufficient understanding of the strengths and limitations of the organisation's AI-enabled systems is vital.

13) Ensure that senior management has a suitable understanding of AI (in relation to their roles and responsibilities).
14) Train risk management and compliance personnel in AI.
15) Develop awareness and understanding of AI within your organisation.

Ad 13: The board of directors' competences and expertise include relevant and up to date knowledge and/or experience enabling the board to sufficiently understand the risks associated with the application of AI to the organisation's core business processes. Managers responsible for AI applications are trained to understand AI-specific challenges, issues and risks.

Ad 14: Risk management and compliance personnel that is responsible for the assessment and audit of AI-enabled systems is trained to understand AI-specific challenges, issues, and risks. Their competence and expertise is of a sufficient level to adequately perform their second line function.

Ad 15: Personnel working with AI applications are trained to responsibly use these applications, and sufficiently understand their strengths and limitations. Educational programs are used to create and improve awareness regarding the use of AI applications and their associated risks throughout the organisation.

4.6 Transparency

Transparency means that financial firms should be able to explain how they use AI in their business processes, and (where reasonably appropriate) how these applications function. Adherence to this principle enables adequate risk management and internal audits. This also means effective supervision of the correct workings of the firm's AI applications to ensure stable operations. When the materiality of AI applications increases, so should the organisation's efforts in understanding and controlling the applications' functioning.

16) Be transparent about your policy and decisions regarding the adoption and use of AI internally.
17) Advance traceability and explainability of AI driven decisions and model outcomes.

Ad 16: The organisation's policy regarding AI is clearly communicated throughout the organisation. Material decisions regarding AI applications, and their underlying models and data, are systematically documented and sufficiently justified. The limitations of adopted models and data sets are well documented and communicated, as well as the circumstances under which the use of a particular model or data set should be discontinued. The reasons for specific model choices for AI applications are systematically documented and available for review. Decisions that favour accuracy over traceability and explainability should be well-motivated, documented, and approved at the appropriate level.

Ad 17: Where reasonable, efforts are made to improve the reproducibility of operations and outcomes of AI systems, in order to facilitate review processes. AI model outcomes are sufficiently traceable and explainable. This means that, where feasible, it can be reasonably demonstrated how (sets of) individual input parameters contribute to the model's outcomes on an aggregate level (global explanation) and/or how individual outcomes tend to respond to changes in the input variables (local explanation). In cases where the relation between input parameters and model outcomes is obscure because of model choices (eg. when deploying advanced convolutional neural networks), domain experts periodically evaluate model outcomes to improve the organisation's understanding of the model's workings. A continuous effort is made to develop processes, tools and interfaces to further improve traceability and explainability of AI applications to ensure continuing alignment of their performance with the organisation's objectives. Procedures are in place for 'after the fact' reviews, on instigation of the customer and/or other stakeholders.

References

ABN AMRO. (n.d.). *Use the Grip app and manage your money*. Retrieved from ABN AMRO: https://www.abnamro.nl/en/personal/apps/grip/index.html

ACPR. (2018). *Artificial Intelligence: challenges for the financial sector.* Paris: Autorité de Contrôle Prudentiel et de Résolution .

Akerlof, G. A. (1970). The market for "lemons": quality uncertainty and the market mechanism. *The Quearterly Journal of Economcs, 84*(3), 488-500.

Artificial Intelligence Applications Institute. (2005). *Edinburgh Freddy Robot (Mid 1960s to 1981).* Retrieved from http://www.aiai.ed.ac.uk/project/freddy/

ASR. (n.d.). *ASR homepage*. Retrieved from https://www.asr.nl/

Bachor, H.-A., & Simmons, M. (2017, November 32). *Quantum computing*. Retrieved from Australian Academy of Science: https://www.science.org.au/curious/technology-future/quantum-computing

BaFin. (2018). *Big data meets artificial intelligence: challenges and implications for the supervision and regulation of financial services.* Frankfurt am Main: Bundesanstalt für Finanzdienstleistungsaufsicht.

Bawden, A., Greenblatt, R., Holloway, J., Knight, T., Moon, D., & Weinreb, D. (1977). *LSIP machine progress report.* Cambridge: Massachusetts Institute of Technology.

Berliner, H. J. (1977). Experiences in evaluation with BKG-A program that plays backgammon. *Proceedings of the 5th international joint conference on Artificial intelligence - Volume 1* (pp. 428-33). San Francisco: Morgan Kaufmann Publishers Inc.

Bouchard, K. E., Aimone, J. B., Chun, M., Dean, T., Denker, M., Diesmann, M., . . . Prabhat. (2018). International Neuroscience Initiatives through the Lens of High-Performance Computing. *Computer, 51*(4), 50-9.

Boudette, N. E. (2017, January 19). Tesla's self-driving system cleared in deadly crash. *New York Times*.

Brown, M. (2017, January 2). *Tesla just launched its futuristic "Enhanced Autopilot" update*. Retrieved from Inverse: https://www.inverse.com/article/25848-elon-musk-confirm-enhanced-autopilot-tesla

Brown, N., & Sandholm, T. (2018, January 26). Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science, 359*(6374), pp. 418-24.

Brunkhard, P. (2018, September 3). *The future of IOT is AI*. Retrieved from TechUK: https://www.techuk.org/insights/opinions/item/13827-the-future-of-iot-is-ai

BusinessForensics. (n.d.). *References*. Retrieved from https://businessforensics.nl/en/references/#clients

Chodosh, S. (2017, April 22). *The incredible evolution of supercomputers' powers, from 1946 to today*. Retrieved from Popular science: https://www.popsci.com/supercomputers-then-and-now/

CSAIL. (n.d.). *Mission & History*. Retrieved from https://www.csail.mit.edu/about/mission-history

CSSF. (2018). *Artificial intelligence: opportunities, risks and recommendations for the financial sector.* Luxembourg: Commission de Surveillance du Secteur Financier .

Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*.

deeplearning.ai. (2018, August 25). *Heroes of Deep Learning: Geoffrey Hinton*. Retrieved from https://www.deeplearning.ai/blog/hodl-geoffrey-hinton/

DeepMind. (2019). *AlphaStar: Mastering the Real-Time Strategy Game StarCraft II*. Retrieved from https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/

Dodds, L. (2018, November 25). Chinese business-woman accused of jaywalking after AI camera spots her face on an advert. *The Telegraph*.

EBA. (2019). *Consultation paper: draft guidelines on loan origination and monitoring.* European Banking Authority.

European Commission. (n.d.). *Artificial intelligence*. Retrieved from European Commission: https://ec.europa.eu/digital-single-market/en/artificial-intelligence

Fairzekering. (n.d.). *Zo werkt onze Autofairzekering*. Retrieved from https://fairzekering.nl/hoe-werkt-het/

FRISS. (n.d.). *FRISS fraud detection for insurance claims and lower loss ratios*. Retrieved from FRISS: https://www.friss.com/products/fraud-detection-at-claims/

FSB. (2017). *Artificial intelligence and machine learning in financial services: Market developments and financial stability implications.* Basel: Financial Stability Board.

Giles, M. (2018, December 12). *Intel's new 3D chip technology may help prolong Moore's Law*. Retrieved from MIT Technology Review: https://www.technologyreview.com/f/612587/intels-new-3d-chip-technology-may-help-prolong-moores-law/

Grand View Research. (2017). *Artificial Intelligence Market Analysis By Solution (Hardware, Software, Services), By Technology (Deep Learning, Machine Learning, Natural Language Processing, Machine Vision), By End-use, By Region, and Segment Forecasts, 2018 - 2025*. Grand View Research.

Greenblatt, R. D., Eastlake, III, D. E., & Crocker, S. D. (1967). The Greenblatt chess program. *Proceeding AFIPS '67 (Fall) Proceedings of the November 14-16, 1967, fall joint computer conference* (pp. 801-10). Anaheim: ACM.

Grush, L. (2015, July 1). *Google engineer apologizes after Photos app tags two black people as gorillas*. Retrieved from The Verge: https://www.theverge.com/2015/7/1/8880363/google-apologizes-photos-app-tags-two-black-people-gorillas

Hern, A. (2015, May 20). Flickr faces complaints over 'offensive' auto-tagging for photos. *The Guardian*.

High-Level Expert Group on Artificial Intelligence. (2019). *Ethics guidelines for trustworthy AI.* Brussels: European Commission.

ING. (2015, October 21). *ING to start strategic partnership and launch pilot with fintech Kabbage*. Retrieved from https://www.ing.com/Newsroom/All-news/Press-releases/ING-to-start-strategic-partnership-and-launch-pilot-with-fintech-Kabbage.htm

ING. (2017, December 12). *Katana gives bond traders a cutting edge*. Retrieved from https://www.ing.com/Newsroom/All-news/Katana-gives-bond-traders-a-cutting-edge.htm

ING. (2018, December 17). *Using AI to assess credit risk*. Retrieved from https://www.ing.com/Newsroom/All-news/Using-AI-to-assess-credit-risk.htm

ING. (2019, March 27). *ING takes academic approach to AI*. Retrieved from https://www.ing.com/Newsroom/All-news/ING-takes-academic-approach-to-AI.htm

JPMorgan Chase & Co. (2017). *Annual report 2016.* JPMorgan Chase & Co.

Kochansky, J. (2018, March 8). *The promise of artificial intelligence and what it means to BlackRock*. Retrieved from BlackRock Blog: https://www.blackrockblog.com/2018/03/08/artificial-intelligence-blackrock/

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*, (pp. 1097-105).

Kubota, T. (2019, January 16). *Stanford's robotics legacy*. Retrieved from Stanford News: https://news.stanford.edu/2019/01/16/stanfords-robotics-legacy/

Le, Q. V., Ranzato, M., Monga, R., Devin, M., Chen, K., Corrado, G. S., . . . Ng, A. Y. (2012). Building high-level features using large scale unsupervised learning. *Proceedings of the 29th International Conference on Machine Learning.* Edinburgh, Scotland, UK.

Linnainmaa, S. (1970). *The representation of the cumulative rounding error of an algorithm as a Taylor expansion of the local rounding errors.* Univ. Helsinki.

Lun, R., & Zhao, W. (2015). A survey of applications and human motion recognition with Microsoft Kinect. *International Journal of Pattern Recognition and Artificial Intelligence*, 29(05).

Lynch, S. (2017, Macrh 14). *Andrew Ng: Why AI is the new electricity*. Retrieved from Stanford News: https://news.stanford.edu/thedish/2017/03/14/andrew-ng-why-ai-is-the-new-electricity/

Madrigal, A. C. (2017, July 19). *How checkers was solved*. Retrieved from The Atlantic: https://www.theatlantic.com/technology/archive/2017/07/marion-tinsley-checkers/534111/

Malecek, A. (2016, September 2). *For first time, carbon nanotube transistors outperform silicon*. Retrieved from Phys.org: https://phys.org/news/2016-09-carbon-nanotube-transistors-outperform-silicon.html

Markoff, J. (2011, February 17). Computer wins on 'Jeopardy!': trivial, it's not. *The New York Times*, p. 1.

Massa, A. (2018, August 30). *BlackRock's Decade: How the Crash Forged a $6.3 Trillion Giant*. Retrieved from Bloomberg: https://www.bloomberg.com/news/articles/2018-08-30/blackrock-s-decade-how-the-crash-forged-a-6-3-trillion-giant

Microsoft. (2011, January 25). *Human pose estimation for Kinect*. Retrieved from Microsoft: https://www.microsoft.com/en-us/research/project/human-pose-estimation-for-kinect/

Milton, S., Carbonell, J. G., Knoblock, C. A., Kuokka, D. R., Etzioni, O., & Gil, Y. (1989). Explanation-based learning: A problem solving perspective. *Artificial Intelligence*, *40*(1-3), 63-118.

Monetary Authority of Singapore. (2018). *Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector.* Singapore: MAS.

Moss, S. (2018, June 8). *The PlayStation Supercomputer*. Retrieved from DatacenterDynamics: https://www.datacenterdynamics.com/analysis/the-playstation-supercomputer/

Mundy, S. (2016, March 15). AlphaGo conquers Korean grandmaster Lee Se-dol. *Financial Times*.

Municipality of Amsterdam. (2018, October 9). *AI report shows Dutch companies are ahead in Europe.* Retrieved from I amsterdam: https://www.iamsterdam.com/en/business/news-and-insights/news/2018/ai-report-shows-dutch-companies-are-ahead-in-europe

Nicolau, D. V., Lard, M., Korten, T., van Delft, F. C., Persson, M., Bengtsson, E., . . . Nicolau, D. V. (2016). Parallel computation with molecular-motor-propelled agents in nanofabricated networks. *PNAS*, *113*(10), 2591-6.

OECD. (1997). *Information Technology Outlook.* Paris: OECD.

Orbis Research. (2017, April 25). *Artificial Intelligence-Global Market Outlook (2016-2022).* Dallas: Orbis Research. Retrieved from Reuters: https://www.reuters.com/brandfeatures/venture-capital/article?id=4999

Oreskovic, A. (2016, March 23). *Intel admits the engine behind the tech industry's amazing 50-year run is slowing down*. Retrieved from Business Insider: https://www.businessinsider.com/intel-acknowledges-slowdown-to-moores-law-2016-3?international=true&r=US&IR=T

Perry, T. S. (2018, April 2). *Move Over, Moore's Law: Make Way for Huang's Law*. Retrieved from IEEE Spectrum: https://spectrum.ieee.org/view-from-the-valley/computing/hardware/move-over-moores-law-make-way-for-huangs-law

Quillian, M. R. (1968). Semantic networks. In M. L. Minsky, *Semantic Information Processing.* MIT Press.

Rapier, G., & Bastone, N. (2019, April 23). *Tesla unveils the 'best chip in the world' for self-driving cars at its autonomy day event*. Retrieved from Business Insider: https://www.businessinsider.nl/tesla-autonomy-investor-day-live-updates-event-details-elon-musk-self-driving-2019-4/

Reinsel, D., Gantz, J., & Rydning, J. (2018). *The digitization of the world: from edge to core.* IDC.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). See Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August). Why should i trust you? Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1135-44). ACM.

Rosenberg, C. R., & Sejnowski, T. J. (1986). The spacing effect on NETtalk, a massively-parallel network. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society* (pp. 72-89). New Jersey: Erlbaum Hillsdale.

Salesforce. (n.d.). *Customer success stories*. Retrieved from https://www.salesforce.com/nl/customer-success-stories/#!page=1

Samuel, A. (1959). Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development*, *3*(3), 210-29.

Sanberg, A., & Bostrom, N. (2008). *Whole brain emulation: a roadmap. Technical report #2008-3.* Future of Humanity Institute, Oxford University.

Sarangi, S., & Sharma, P. (2019). *Artificial Intelligence: evolution, ethics and public policy.* New York: Routledge. Retrieved from https://www.dartmouth.edu/~ai50/homepage.html

Shortliffe, E. H. (1977). Mycin: A Knowledge-Based Computer Program Applied to Infectious Diseases. *Proceedings of the Annual Symposium on Computer Application in Medical Care 1977*, (pp. 66-9).

Sicular, S., & Brant, K. (2018). *Hype cycle for Artificial Intelligence, 2018.* Gartner, Inc.

Silver, D., Hubert, T., Schrittweiser, J., Antonoglou, I., Lai, M., Guez, A., . . . Hassabis, D. (2018, December 7). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science, 362*(6419), pp. 1140-4.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., . . . Hassabis, D. (2017, October 19). Mastering the game of Go without human knowledge. *Nature*, pp. 354-359.

Solomonoff, R. J. (1960). *A preliminary report on a general theory of inductive inference.*

Sonnad, N. (2017, November 29). *Google Translate's gender bias pairs "he" with "hardworking" and "she" with lazy, and other examples*. Retrieved from Quartz: https://qz.com/1141122/google-translates-gender-bias-pairs-he-with-hardworking-and-she-with-lazy-and-other-examples/

Stanford Research Institute. (n.d.). *Shakey the Robot*. Retrieved from SRI International: https://www.sri.com/work/timeline-innovation/timeline.php?timeline=computing-digital#!&innovation=shakey-the-robot

Statista. (2015, December 16). *Robotics and artificial intelligence (AI) worldwide market size estimates, based on 2018 to 2030 forecasts, by segment (in billion U.S. dollars)*. Retrieved from Statista: https://www.statista.com/statistics/621656/worldwide-artificial-intelligence-robotics-segment-estimates/

Statista. (2016, September 30). *Share of companies investing in artificial intelligence (AI) worldwide, by industry, as of 2016*. Retrieved from Statista: https://www.statista.com/statistics/621544/worldwide-artificial-intelligence-investment-by-industry/

Statista. (2019, April 5). *Share of big data and business analytics revenues worldwide in 2019, by industry*. Retrieved from Statista: https://www.statista.com/statistics/616225/worldwide-big-data-business-analytics-revenue/

TOP500.org. (2019, June). *TOP500 List*. Retrieved from https://www.top500.org/lists/2019/06/

Tractica. (2019). *Artificial Intelligence Market Forecasts.* Tractica.

Turing, A. (1950). Computer machinery and intelligence. *Mind*, 49(236), 433-60.

Uhr, L., & Vossler, C. (1961). A pattern recognition program that generates, evaluates, and adjusts its own operators. *Papers presented at the May 9-11, western joint IRE-AIEE-ACM computer conference.* New York: ACM.

Vincent, J. (2016, March 24). *Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day*. Retrieved from The Verge: https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist

Vincent, J. (2019, May 7). *World's fastest supercomputer will be built by AMD and Cray for US government.* Retrieved from The Verge: https://www.theverge.com/2019/5/7/18535078/worlds-fastest-exascale-supercomputer-frontier-amd-cray-doe-oak-ridge-national-laboratory

Weber, B. (1997, May 12). Swift and slashing, computer topples Kasparov. *The New York Times*, p. 1.

WEF. (2018). *The New Physics of Financial Services: Understanding how artificial intelligence is transforming the financial ecosystem.* World Economic Forum.

WEF. (2019). *Fourth Industrial Revolution*. Retrieved from World Economic Forum: https://www.weforum.org/focus/fourth-industrial-revolution

Wegdam, M. (2018, February 4). *Rabobank starts mobile identification*. Retrieved from InnoValor: https://innovalor.nl/en/Blogs/Rabobank-starts-with-mobile-identication

Weiser, M. (1991). The computer for the 21st century. *Scientific American*, 265(3), 94-105.

Weizenbaum, J. (1966). ELIZA---a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36-45.

Welch, D., & Behrmann, E. (2018, May 7). *Who's Winning the Self-Driving Car Race?* Retrieved from Bloomberg: https://www.bloomberg.com/news/features/2018-05-07/who-s-winning-the-self-driving-car-race

Winograd, T. (1971). *Procedures as a representation for data in a computer program for understanding natural language* . Cambridge: Massachusetts Institute of Technology.

World Wide Web Foundation. (n.d.). *History of the Web*. Retrieved from https://webfoundation.org/about/vision/history-of-the-web/

Yadav, N., Yadav, A., & Kumar, M. (2015). *An introduction to neural network methods for differential equations.* Dordrecht: Springer.

Appendix A Relevant articles in the Financial supervision Act (unofficial translation)

Financial Supervision Act section 3:10

1. A clearing institution, credit institution or insurer having its registered office in the Netherlands shall pursue an adequate policy that safeguards controlled and sound business operations. This shall mean that:
   a. measures are taken to prevent conflicts of interest;
   b. measures are taken to prevent the financial enterprise or its employees from committing offences or other transgressions of the law that could damage confidence in the financial enterprise or in the financial markets;
   c. measures are taken to prevent confidence in the financial enterprise or in the financial markets from being damaged because of its clients; and
   d. measures are taken to prevent the financial enterprise or its employees from performing other acts that are so contrary to generally accepted standards as to seriously damage confidence in the financial enterprise or in the financial markets.

Financial Supervision Act section 3:17

1. A clearing institution, credit institution or insurer having its registered office in the Netherlands shall organise its operations in such a way as to safeguard controlled and sound business operations.
2. Rules shall be laid down by or pursuant to a Decree with regard to Subsection (1). These rules shall concern:
   a. control of business processes and business risks;
   b. integrity, which is understood to mean the prevention of:
      i. conflicts of interest;
      ii. offences or other transgressions of the law committed by the financial enterprise or its employees that could damage confidence in the financial enterprise or in the financial markets;
      iii. relations with clients that could damage confidence in the financial enterprise or in the financial markets; and
      iv. other acts performed by the financial enterprise or its employees that are so contrary to generally accepted standards as to seriously damage confidence in the financial enterprise or in the financial markets;
   c. the soundness of the financial enterprise, which shall be understood to mean:
      i. control of financial risks;
      ii. control of other risks that may affect the soundness of the financial enterprise;
      iii. ensuring the maintenance of the required financial safeguards; and
      iv. other matters, to be specified by Decree.
3. Without prejudice to Section 4:14, Subsection (2), opening words and under (c) shall apply mutatis mutandis to management companies

of a collective investment scheme having its registered office in the Netherlands that offers units in the Netherlands, collective investment schemes having their registered office in the Netherlands that offer units in the Netherlands, investment firms having their registered office in the Netherlands that provide investment services or perform investment activities in the Netherlands, and depositaries associated with a collective investment scheme having its registered office in the Netherlands that offers units in the Netherlands.

Financial Supervision Act section 3:18

1. If a financial enterprise having its registered office in the Netherlands delegates activities to a third party, the financial enterprise shall ensure that such third party complies with the rules applicable under this part to the delegating financial enterprise with regard to those activities.
2. A clearing institution, credit institution or insurer shall not delegate activities to be designated by Decree.
3. By or pursuant to a Decree:
   a. rules shall be laid down, for the purpose of supervision of compliance with the provisions arising from this part, with regard to the delegation of activities by financial enterprises;
   b. rules shall be laid down with regard to the control of risks entailed by the delegation of activities by clearing institutions, credit institutions and insurers; and
   c. rules shall be laid down with regard to the agreement to be concluded between a clearing institution, credit institution or insurer and the third party relating to the delegation of activities.

DeNederlandscheBank

EUROSYSTEEM

De Nederlandsche Bank N.V.
PO Box 98, 1000 AB Amsterdam
+31 20 524 91 11
dnb.nl