

**The
Alan Turing
Institute**

**Artificial
intelligence
in finance**

Bonnie G. Buchanan, PhD, FRSA



The Alan Turing Institute

<https://doi.org/10.5281/zenodo.2612537>

This work was supported by The Alan Turing Institute
under the EPSRC grant EP/N510129/1

Artificial intelligence in finance

April 2019

Bonnie G. Buchanan, PhD, FRSA

Howard Bosanko Professor of Economics and Finance

Department of Finance,

Albers School of Business and Economics

Seattle University

Seattle, Washington 98122-1090

Email: buchanab@seattleu.edu

Ph: (206) 296-5977

Hanken School of Economics

Department of Finance, Statistics and Economics

P.O. Box 479, FI-00101 Helsinki, Finland

Abstract

Artificial intelligence (AI) is rapidly transforming the global financial services industry. As a group of related technologies that include machine learning (ML) and deep learning (DL), AI has the potential to disrupt and refine the existing financial services industry. I review the extant academic, practitioner and policy related AI literature. I also detail the AI, ML and DL taxonomy as well as their various applications in the financial services industry.

A literature survey of AI and financial services cannot ignore the econometric aspects and their implications. ML methods are all about algorithms, rather than asymptotic statistical processes. Unlike maximum likelihood estimation, ML's framework is less unified. To that end, I will discuss the ML approaches of unsupervised and supervised learning.

Contents

1. Introduction	1
2. Taxonomy and historical overview of AI, ML and DL	4
3. Global growth of the AI industry	7
4. How AI is changing the financial services industry	11
4.1. Fraud detection and compliance	11
4.2. Banking chatbots and robo-advisory services	13
4.3. Algorithmic trading	15
4.4. Other applications of AI.....	18
5. Econometrics versus ML	19
5.1. Unsupervised machine learning	20
5.1.1. Clustering algorithms.....	20
5.1.2. Topic models	20
5.1.2. Topic models (cont.).....	22
5.2. Supervised machine learning models	22
5.2.1. Predictive analytics.....	23
5.2.2. Random forests.....	23
5.2.3. Neural networks	23
5.2.4. Support vector machine (SVM).....	24
6. Machine learning versus quantum computing	25
7. Regulation and policy-making	26
8. Conclusion and directions for future research	29
Appendix	30
Timeline of artificial intelligence milestones.....	30
References	32

“AI is the ‘new electricity’ ... just as electricity transformed many industries roughly one hundred years ago; AI will also now change every major industry.”

Andrew Ng, 2007

“What we're seeing is something unprecedented, which is the arrival of artificial intelligence, which has a big impact ... it creates tremendous uncertainty and impacts different people differently ... and some people could be left out.”

Robert Shiller, 2018 Davos Forum

1. Introduction

In 1950, Alan Turing posed the question “Can machines think?” and since then artificial intelligence (hereafter known as AI) applications have met with varying degrees of success. However, in recent years there has been a resurgence of interest and AI has found innovative applications in the global financial services industry. The availability of big data, improved technology, cloud computing and faster special purpose hardware have been key drivers of the latest AI innovation wave. AI capabilities and machine learning (ML) are boosting growth in the emerging Fintech market. Broadly speaking, the term “Fintech” describes the new technologies, services and companies that have changed financial services. It includes (but is not limited to): cryptocurrencies, blockchain¹, robo-advising, smart contracts, crowdfunding, mobile payments and AI platforms. In 2017 AI topped the list as a key trend in financial services and Fintech (Future Today Institute, 2017).

In this literature review, I will detail the AI, ML and deep learning (DL) taxonomy as well as their various applications in the financial services industry. I will summarise the current academic, practitioner and policy related AI literature. This includes drawing upon economic, finance and computer science literature as well as regulatory publications. I specifically discuss four ways in which AI is changing the financial services industry: (1) fraud detection (how AI is used to keep criminal funds out of the financial system); (2) banking chatbots; (3) algorithmic trading and (4) regulatory and policy aspects.

Professor John McCarthy coined the term “artificial intelligence” in 1955 and the term “machine learning” was coined in 1959 by Arthur Samuel of IBM. AI can mimic actions it has seen or previously have been taught about without any new intervention. ML is defined as a particular approach to AI able to take the data and algorithms and apply it to new scenarios and patterns without being programmed directly. Deep learning (DL) is viewed as a branch of ML. DL provides machines with algorithms necessary to understand the underlying principles of an action and significant portions of data. They can then be combined to learn on their own and deepen the knowledge and skills with which they are provided. High frequency trading (HFT) and algorithmic trading use high speed communications and computer programs in the financial services industry. For at least a decade banks have been using ML to detect credit card fraud.

The UK Financial Conduct Authority (FCA) is utilising ML to help individuals manage their current accounts. Approximately 9% of all hedge funds use ML to build large statistical models. In 2016, Aidya launched an AI hedge fund to make all its stock trades. Sentient Investment Technologies uses a distributed AI system and DL as part of its trading and investment platform. Fukoku Mutual Life Insurance uses IBM’s Watson Explorer AI to calculate pay-outs. Feedzai uses ML to detect fraudulent transactions. UK PropTech² start-up Leverton applies AI to automatically identify, extract and manage data from corporate documents such as rental leases. In October 2017, exchange traded funds (ETFs) were launched that use AI algorithms to choose long-term stock holdings.

Like other Fintech sectors, AI offers many opportunities and challenges. In terms of financial inclusion, the increased application of AI technology to capital markets is likely to reduce barriers to entry for many individuals who might not have previously had access to financial markets. Some of the world’s most valuable big tech companies such as Apple, Amazon, Tencent and Alibaba have been pouring money into AI research. But as Robert Shiller’s remarks at the 2018 Davos Forum indicate, AI also presents a great deal of uncertainty as a disruptive technology.

In 2016, the GIS-Liquid Strategies group was managing \$13 billion with only 12 people. In 2017 Standard & Poor’s (S&P) acquired Kensho for \$550 million in the biggest AI acquisition to date. Kensho was founded in 2013 with the intention of replacing bond and equity analysts. Its algorithm is dubbed “Warren” (after Warren Buffet) and it can process 65 million question combinations by scanning over 90,000 events such as economic reports, drug approvals, monetary policy changes and political events and their impact on financial assets. DeepMind Technologies was purchased by Google and Intel has acquired Nervana Systems. In 2017 Opimas LLC estimated that AI would result in approximately 230,000 job cuts in financial firms worldwide by 2025, with the hardest hit area being asset management (with an estimated 90,000 job cuts)³.

A literature survey of AI and financial services cannot ignore the econometric aspects and implications. ML methods are about algorithms, more than about asymptotic statistical processes⁴. Unlike maximum likelihood estimation, ML’s framework is less unified. To that end, I will discuss the ML approaches of unsupervised and supervised learning.

In unsupervised learning, ML can help issue account alerts such as low balance warnings. It can also be applied to bank overdraft charges to help ascertain what is happening to individual customers and what might be the causes of the situation. This is accomplished by using clustering algorithms. Regulators can also use clustering algorithms to better understand trades and categorise business models of banks in advance of regulatory examinations.

The SEC is using topic models to detect accounting fraud. Topic models help us understand the behavioural drivers of different market participants. Topic models draw on text mining and natural language processing (NLP). Both of these unsupervised techniques are precursors to predictive analytics (or supervised ML). Supervised ML entails teaching an algorithm to learn from past breaches of regulations and predict new breaches, insider trading and cartel detection. In this literature survey I will also discuss Random Forests⁵, neural networks (a type of deep learning), as well as least absolute selection and shrinkage operator (LASSO) regressions.

Additionally, I plan to address the following knowledge gap - ML is anticipated to have a far greater potential impact if it is combined with the processing capabilities of quantum computing. If quantum computing becomes a reality, it has the potential to disrupt blockchain. Once this is achieved, what does this mean for the financial services industry?

AI continues to become more sophisticated and complex, but so do the financial markets and this presents major challenges in regard to regulation and policy-making. Finally, I will discuss how ML is making an impact on the tools regulators use to set policy, detect fraud, estimate supply and demand and ensure compliance. In the future, regulators will still need to have procedures in place for determining whether a firm or person is at fault. I will detail the international regulatory responses to AI and financial services, with an emphasis on the UK. I will consider examining various UK agencies such as the Bank of England, the FCA, Serious Fraud Office and the Competition and Markets Authority.

The paper is set out as follows: Section Two describes the taxonomy and historical overview of AI, ML and DL. Section Three details the global growth of AI, followed by three examples of how AI is changing the financial services industry in Section Four. In Section Five I describe at length the differences between various ML techniques and traditional econometric methods. The impact of the emerging field of quantum computing on AI is discussed in Section Six. The regulatory response to AI is provided in Section Seven and Section Eight concludes.

2. Taxonomy and historical overview of AI, ML and DL

The term “artificial intelligence” was coined in 1956 by John McCarthy. The Oxford English Dictionary defines AI as “The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making and translation between languages.”⁶ FSB (2017) defines AI as, “The theory and development of computer systems able to perform tasks that have traditionally required human intelligence.”

Both are fairly broad definitions. Kaplan (2016) describes AI as, “The essence of AI, indeed the essence of intelligence, is the ability to make appropriate generalizations in a timely fashion based on limited data. The broader the domain of application, the quicker the conclusions are drawn with minimal information, the more intelligent the behavior.”

The historical development of the AI field is presented in the **Appendix**. It is clear from the timeline that early AI efforts concentrated on rules that included logic-based algorithms. Turing (1950) details an operational test (the Turing Test) for intelligent behaviour. In his seminal work, Turing provided the major components for future AI work with language, reasoning, knowledge, learning and understanding. Through the Turing Test, Turing laid the ground work for ML, genetic algorithms and reinforcement learning. The attempt to replicate the logical flow of human decision making through processing symbols became known as the “symbol processing hypothesis” (Newell, Shaw, and Simon, 1957; Newell and Simon, 1961, Gilmarin, Newell and Simon, 1976).

Much of AI in the 1950s and 1960s did not focus on finance applications. In the 1960s, a substantial body of work on Bayesian statistics was being developed that would later be used in ML. Neural networks (which would become a cornerstone of deep learning) were developed in the 1960s and grew rapidly. However, due to a lack of sufficiently available electronic data and computing power, AI fell out of favour into what became known as an “AI winter” (Kaplan, 2016; FSB, 2017). The term “AI Winter” also connotes a slowdown in investment and interest. In 1973, the UK Lighthill Report ended government support for AI research.

The 1980s witnessed an AI revival due to new funding and techniques. During the 1980s, Japan, the UK and the USA competed heavily in AI funding. Japan invested \$400 million through the Japanese Fifth Generation Computer Project. The UK invested £350 million in the Alvey Program and DARPA spent over \$1 billion on its Strategic Computing Initiative. In 1982 AI made inroads into the financial services industry when James Simons founded quantitative investment firm Renaissance Technologies⁷.

This included the development of “expert systems” (or “knowledge systems”) which is a technique that solves problems and answers questions within a specific context. Brown, Nielson and Phillips (1990) provide an overview of integrated personal financial planning expert systems. They emphasise expert systems that use heuristics and the separation of knowledge and control as well as providing examples of expert systems that were prevalent at the time. For example, PlanPower provided tailored financial plans to individuals with incomes over \$75,000.

The Personal Financial Planning System (PFPS) was used by Chase Lincoln First Bank and Arthur D. Little Inc. to undertake investment planning, debt planning, retirement planning, education planning, life-insurance planning, budget recommendations and income tax planning.

Expert systems were also used in the stock market in what was known as “program trading.” At the time, institutional investors used program trading to capitalise on pricing disparities in the market. Finnerty and Park (1987) provide an empirical study of program trading that identifies discrepancies between stock index futures and the underlying stock index. They find the program trading strategy consistently outperforms the simple execute and hold to expiration strategy.

However, program trading was often attributed to the wild market swings of the late 1980s, culminating in the 508-point drop in the Dow Jones Industrial Average (DJIA) in 1987. Chen and Liang (1989) detail PROTRADER (an expert system prototype for program trading) which is based on a learning mechanism based on parameter adjustment to several critical parameters based on market conditions. Chen and Liang (1989) describe a successful prediction of the DJIA 87 point drop in 1986.

Even though expert systems have since declined, one notable example that still remains is Fair Isaac Corporation's (FICO)⁸ Blaze Advisor business rules management system (Kaplan, 2016). The late 1980s witnessed the rise of IBM and Apple desktop computers. As specialised expert systems became more expensive to maintain, a second AI winter ensued. This was also driven by companies that had failed to deliver on extravagant promises. This second AI winter lasted until approximately 1993. Despite this in 1988, David Shaw founded a hedge fund (D.E. Shaw) that was an early adopter of AI techniques for trading.

After the work of Cheeseman (1985) and Pearl (1988) Bayesian analysis gained greater acceptance in addressing uncertainty in AI research. The 1960s and 1970s had been dominated by probabilistic reasoning models but Bayesian networks combined classical AI and neural nets and allowed for learning from experience.

In the 1990s, the use of AI in fraud detection garnered more interest. The FinCEN Artificial Intelligence System (FAIS)⁹ was put into service in 1993 in an effort to predict and assess money laundering incidents. Over the following two years, FAIS would review over 200,000 transactions per week and identify 400 potential money laundering incidents, worth approximately \$1 billion.

After the second AI winter concluded, AI advanced more into new areas like machine learning, data mining, virtual reality and case-based reasoning. ML is considered a subset of AI and uses algorithms to automatically optimise through experience with limited or no human intervention. ML is primarily derived from sources such as experience, practice, training and reasoning. More specifically, ML is concerned with general pattern recognition and universal approximations of relations in data in cases where no a priori analytical solution exists (Cybenko, 1989).

Machine learning is acknowledged to have originated with the work of McCulloch and Pitts (1943). They recognised that brain signals are digital in nature, more specifically binary signals. According to Chakraborty and Joseph (2017) each ML system comprises five components: (1) a problem, (2) data source, (3) a model, (4) an optimization algorithm and (5) validation and testing.

ML is best suited for situations that require extracting patterns from noisy data or sensory perception—or a data-up approach. The four main drivers of the growth in ML include: the transition from physical to electronically stored data; improvements in memory and computing speed, easier access due to the internet; and low-cost high-resolution digital sensors.

By 2011, computer processing power and storage were growing at an accelerating rate and supercomputers presented new possibilities for the AI field. Deep Learning became the new AI breakthrough. FSB (2017) defines “Deep learning” (DL) as “...a form of machine learning that uses algorithms that work in ‘layers’ inspired by the structure and function of the brain. DL algorithms, whose structure are called artificial neural networks, can be used for supervised, unsupervised, or reinforcement learning.” DL is a statistical technique for finding patterns in large amounts of data. It uses neural networks which are based on mimicking the way multiple layers of the brain’s *neurons* work (hence the term “deep”).

Artificial neural networks were first introduced by McCulloch and Pitts (1943) and are architectures which are designed to mimic the way the human brain works. The two primary advantages of DL are: (1) it is more resilient than ML to overfitting and (2) DL can address non-linear events such as market volatility, which usually have to be adjusted manually in standard quantitative models. In 2012 Google had two DL projects in place and by 2017 it had over 1,000 DL projects underway (Makridakis, 2017).

In the last few years, verification in AI has proved to be essential. On August 1, 2012 Knight Capital lost \$440 million in 45 minutes after deploying unverified trading software. The “Flash Crash” on 6 May 2010 was noteworthy for another reason. P&G swung in price between a penny and \$100,000 the problem wasn’t caused by bugs or computer malfunctions that verification could have avoided. It was caused by expectations being violated: automatic trading programs from many companies found themselves operating in an unexpected situation where their assumptions were not valid.

In 2013, during a 17-minute computer glitch, Goldman Sachs flooded the US market with orders to purchase 800,000 contracts linked to equities and ETFs. During the same week, Chinese brokerage firm Everbright Securities, suffered a malfunction which resulted in it purchasing nearly \$4 billion worth of shares on the Shanghai market.¹⁰ After the Brexit referendum in June 2016, Betterment LLC (a new robo-advisor) suspended trading in response to market volatility in order to spare its clients higher transactions costs. Betterment LLC relied heavily on algorithmic trading and regulators were concerned about when and how Betterment LLC might restrict trading (Brummer and Yadav, 2019).

3. Global growth of the AI industry

In the last 60 years the AI field has experienced its share of successes and failures. Currently, governments around the world are competing to create superior AI facilities and research with a view to AI being a lever for greater economic power and influence. Between 2012 and 2016 the US invested \$18.2 billion into AI compared with \$2.6 billion in China and \$850 million in the UK¹¹. The Japanese Government Pension Investment Fund (the world's biggest manager of retirement savings) is considering AI to ultimately replace human fund managers. In February 2018, BlackRock announced it would establish an AI lab¹². With \$6.3 trillion assets under management, the firm already employs text analysis and analyses corporate website traffic and smartphone geolocation data and is now looking at ML to deploy in asset management.

However, the recent trend has been one of rapid growth. According to a Wushen Institute Report (2017), 5,154 AI startups have been established globally during the past five years, representing a 175% increase relative to the previous 12 years. There are two explanations for this impressive growth. First, exponential advances in computing power have led to declining processing and data storage costs and secondly, the immense data availability has increased, creating more possibilities in the AI field.

Historically, the US has dominated the AI industry. Between 2000 and 2016 there were 3,033 AI startups in the US, accounting for 37.41% of the worldwide total (Buchanan and Cao, 2018). However, the proportion has been decreasing and in 2016 dropped to under 30% for the first time. During the same period, the US received \$20.7 billion in funding, accounting for 71.78% of the world's total funding (Wushen Institute Report, 2017)¹³.

In 2017 China surpassed the US for the first time in terms of AI startup funding (CB Insights, 2018). In 2012 China accounted for 48% of global AI startup funding and in 2017 the total global AI funding was \$15.2 billion. AI equity deals increased 141% relative to the previous year and since 2016 more than 1,100 new AI companies have raised their first round of equity financing. However, the US is losing its global AI equity deal share, decreasing from 77% to 50% of equity deal share during the last five years (CB Insights, 2018).

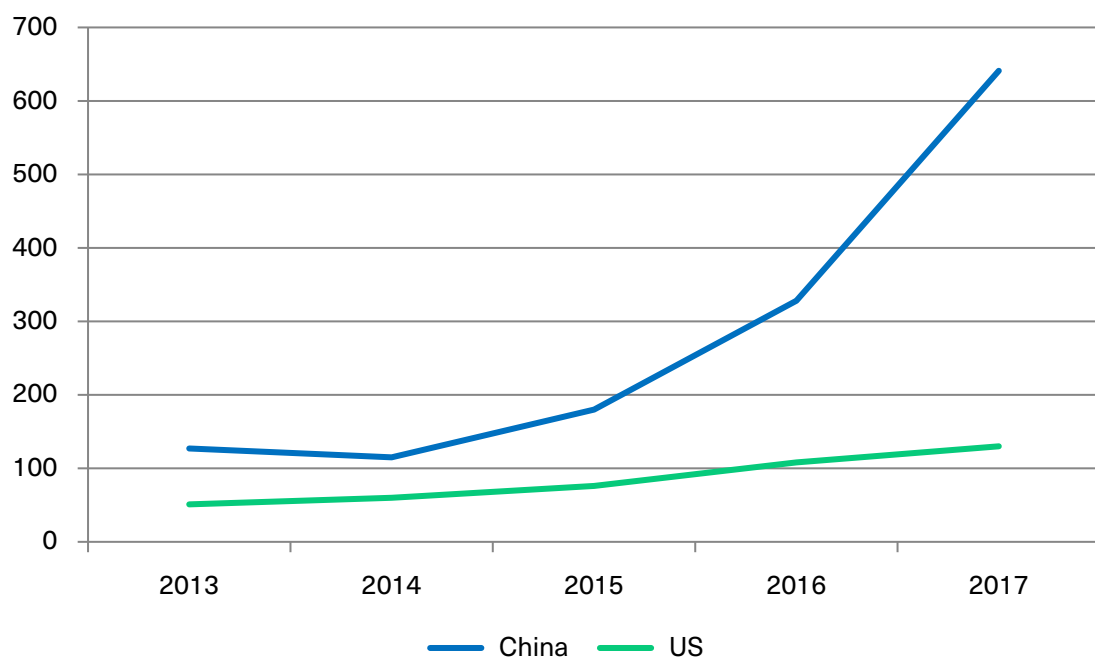
In terms of AI growth, China leads the Asian market. During the past five years China accounted for 68.67% of Asian AI startups and corresponding AI funding was 60.22% of the Asian total. Many Chinese cities and provinces dominate other Asian countries. In terms of the number of AI companies, there are 454 in Beijing, 319 in Guangdong and 224 in Shanghai compared with 57 in Singapore and 283 in India (Wushen Institute Report, 2017).

Beijing has attracted \$1,387 million in AI funding, followed by Guangdong (\$792 million) and Shanghai (\$154 million), the total exceeding that of Japan (\$436.81 million). Beijing's AI current funding is also higher than the UK (\$1,251 million) (Wushen Institute Report, 2017). Within China growth has been swift in the last three years and Beijing accounts for the majority of AI funding (50.23%), followed by Guangdong (28.68%) and Shanghai (5.57%) respectively.

China has overtaken the US for the number of AI patents over the last five years (Buchanan and Cao, 2018). Between 2012 and 2016 the compound growth rate of AI patents was 33.2% per annum. Currently the US and China hold over 50% of all AI patents (35,508 in US and 34,345 in China) but they are being filed at a faster pace in China. The US owns 32% of ML patents and 26% of natural language processing patents around the globe. In these same areas, China comes second with 23% and 14% of patents respectively¹⁴.

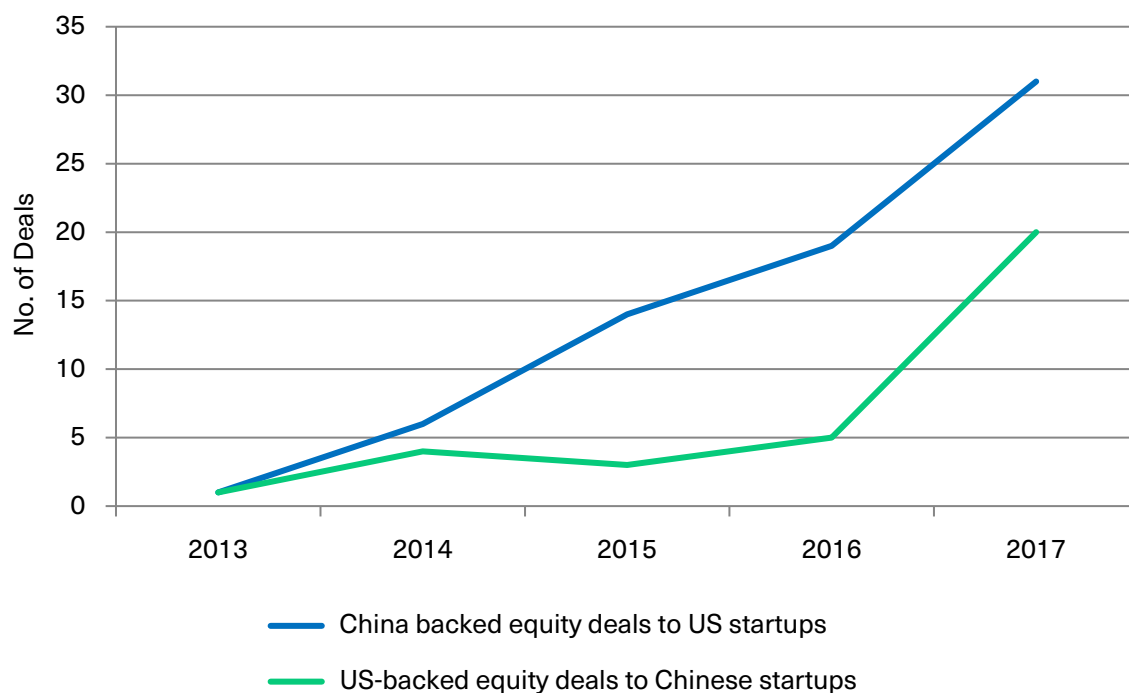
Currently China dominates the machine vision patent category (55% of 150,000 patents globally). Machine vision refers to object and facial recognition and is useful in public security, healthcare, e-commerce and autonomous driving. Two Chinese companies—Face# and SenseTime—are valued as “unicorns.” SenseTime is the most well-funded AI startup to date, with a valuation of US\$4.5 billion, making it the highest valued AI company in history. As **Figure 1** (below) indicates, China exceeds the US in terms of AI related patent publications (by a factor of more than 5 times) between 2013 and 2017.

Figure 1 – AI related patent publications, keyword ‘artificial intelligence’. Source: CB Insights, 2018.



For ML related patent publications, the gap is closing between China and the US. Another useful measure is paper submissions at conferences. In 1980, there was not a single paper from China submitted to the Association for the Advancement of AI Conference. However, in 2018, China submitted 25% more papers than US authors.¹⁵ Additionally, cross border investment is on the rise, but not in an equal manner. As **Figure 2** (next page) indicates there are more Chinese investments in AI startups in the US, than vice-versa. In fact, the gap between the two widened after 2015.

Figure 2 – Cross-border AI investments are on the rise. Source: CB Insights, 2018



So what explains the increasing dominance of China in the AI area? First, there is the matter of scale and China has a massive supply of data. China's online population (730 million) is almost twice the size of the US. Many ML techniques require vast amounts of data. WeChat has over a billion users. Secondly, there are two prominent technologies fuelling the drive, namely facial recognition and AI chips.¹⁶

Many applications like deep neural networks need to supplement central processing units (CPU) with other processors. AI chips draw on graphic processing units (GPU) technology which is then applied to AI, ML or DL problems.¹⁷ Both the US and China compete heavily in AI chip technology. Alibaba is aiming to have its first AI chips on the market in 2019. Currently, most chips are coming from US chipmakers like Qualcomm and Nvidia.

In addition, Google has created the Tensor Processing Unit (TPU) and Microsoft uses a processor called the field programmable gate array (FPGA). Alibaba is aiming to have its first AI chips on the market in 2019. Baidu and JD.com invest heavily in AI both domestically and abroad. Chinese firm, Horizon Robotics is also entering this area. Backed by Intel, Horizon Robotics raised up to \$1 billion in a funding round. Horizon's AI chips run facial recognition algorithms and identify faces from a database of up to 50,000 faces.¹⁸ The AI chips for Horizon's self-driving cars and robots are based on chips that run pre-trained DL algorithms.¹⁹

Chinese companies like Tencent, Baidu and JD.com are Fintech leaders in the Chinese market in terms of market capitalisation and user numbers. All three companies are investing heavily in AI.

However, Chinese and American big tech firms have differed in terms of their AI focus. Microsoft, Google and IBM focus on ML, speech recognition and speech synthesis whereas Tencent, Alibaba and Baidu focus on image recognition and AI searching.²⁰ Chinese company Cambricon is developing chips for DL. Ant Financial uses facial recognition for payments at Alibaba owned retail stores. In 2016, Ant Financial, Foxconn and the city of Hangzhou partnered for the “City Brain” project using AI data from social feeds and surveillance cameras. Additionally, 55 cities participate in the “sharp eyes” project whose surveillance data may end up powering the nation’s Social Credit System, a measure to gauge “trustworthiness” (CB Insights, 2018).

Chinese local governments are offering financial incentives to encourage AI related innovations. Baidu, Alibaba and Tencent also collect vast amounts of data on what consumers buy, where they travel, and who they chat to online. In fact, Baidu has shifted its business strategy from “mobile first” to “AI first.” Baidu’s IPO float of iQiyi is anticipated to provide more additional resources to invest in AI ventures.

In July 2017 the China State Council announced a development plan with the goal of being the world leader in AI by 2030. The council presented a timeline where it expects companies and research facilities to be at parity with the US by 2020.²¹ In June 2017 the Chinese government announced plans to set up an “intelligence industry zone” near Tianjin along with a \$5 billion fund to support the AI industry.

In terms of AI in the United Kingdom, a 2017 government report estimates that AI could add an additional USD \$814 billion (£630 billion) to the UK economy by 2035 (Hall and Pesenti, 2017). In the UK a new AI startup has been established on roughly a weekly basis since 2014.²² Swiftkey, DeepMind and Ravn are among the more prominent AI companies with London being the main hub of AI startups. Bristol company Graphcore is a leading innovator in the chip market. In 2015 BlackRock acquired automated investment platform FutureAdvisor and invested £60 million into early-stage UK venture capital fund Forward Partners, which targeted AI investment.²³ Hall and Pesenti (2017) identify three areas of finance in the UK where AI has great potential: personalised financial planning; fraud detection and anti-money laundering; and process automation.

"Finance is really a wonderfully pure information-processing business."

D.E. Shaw, 1996²⁴

4. How AI is changing the financial services industry

Outside of the technology sector, the financial services industry is the biggest spender on AI services and is experiencing very fast growth (Citi, 2018). Until recently hedge funds and HFT firms were the main users of AI in finance, but applications have now spread to other areas including banks, regulators, Fintech, insurance firms to name a few.

Within the financial services industry, AI applications include algorithmic trading, portfolio composition and optimisation, model validation, back testing, robo-advising, virtual customer assistants, market impact analysis, regulatory compliance and stress testing. In this section, I discuss three specific areas in which AI is currently changing the financial services industry, namely (1) fraud detection and compliance; (2) banking chatbots and robo-advisory services; and (3) algorithmic trading.

4.1. Fraud detection and compliance

As e-commerce has become more widespread, online fraud has also increased. According to the FCA, UK banks spend £5 billion a year in combating financial crime.²⁵ Action Fraud reports that between 2015 and 2016 there was a 66% increase in the number of reported cases of payments-related fraud²⁶ in the UK. US banks spend over \$70 billion on compliance each year. Many large banks have had massive fines imposed upon them for failing to stop illegal financing, and as a result, many banks have turned to AI techniques to improve their operations.

"Benford's Law" is one of the simplest ways to detect fraud. It is accomplished by running an analysis on the first digits in a given set of data. A predictable distribution of first digits will exist in a set of "real" data. Benford's Law has existed since the late 1800s.

AI is beneficial here because ML algorithms can analyse millions of data points to detect fraudulent transactions that would tend to go unnoticed by humans. At the same time ML helps improve the precision of real-time approvals and reduces the number of false rejections. Fraud detection now involves more than a checklist of risk factors. Using ML techniques, fraud detection systems can now actively learn and calibrate in response to new potential (or real) security threats. Using ML, banks' systems can detect unique activities or behaviours ("anomalies") and flag them for investigation.

Credit card fraud detection is one of the most successful applications of ML. Banks are equipped with monitoring systems, or workflow engines, that are trained on historical payments data. Algorithm training, back testing and validation are based on very large datasets of credit card transaction data. Classification algorithms can label events as "fraud" versus "non fraud" and fraudulent transactions can then be stopped in real time (van Liebergen, 2017).

Many financial services companies are exploring AI-based fraud prevention alternatives. Mastercard launched Decision Intelligence (DI) technology and its approach is relatively straightforward. Rather than staying limited to a predefined set of rules, DI derives patterns from historical shopping and spending behaviour of Mastercard customers to establish a baseline. It is against this baseline that Mastercard compares and scores each new customer transaction. This is a major improvement over traditional fraud prevention technologies, which rely on a generic template to evaluate all transactions.

NatWest has adopted a ML solution, namely Corporate Fraud Insights from Vocalink Analytics to detect and prevent redirection fraud. Redirection fraud occurs when a business is tricked into paying money into a fraudster's account rather than their intended supplier.²⁷ In fighting money laundering HSBC uses Quantexa AI software, a UK based start-up.²⁸

Fraud is a latent variable, meaning it is not directly observable but must be inferred from data. As a result, it is more challenging for ML algorithms to make accurate predictions of possible fraud than shopping decisions (a manifest variable), where retailers have access to full transaction histories (Baugess, 2017). Additional challenges exist. A self-defeating goal can arise from declining transactions too aggressively in order to prevent fraud.

A Javelin Strategy Report (2015) indicates that transactions wrongly declined due to suspected fraud — known as a “false positive” — may represent just as big a threat to the financial services industry. This works against the issuer because a false-positive declined transaction can result in erosion of customer loyalty. The same report finds that these “false positives” account for \$118 billion in retail losses and nearly 39% of declined cardholders report that they abandoned their card after being falsely declined.

“Transactions wrongly declined due to suspected fraud account for \$118 billion in retail losses.”

Javelin Strategy Report, 2015

As a response, new technology such as EMV (an acronym for Europay, Mastercard Visa and is also known as smart or chip cards) and mobile payments has helped improve card authorisation practices and reduce false-positive rates (Javelin, 2015). Lopez de Prado and Lewis (2018) examine why false positives are so prevalent in the financial services industry. Citi (2018) estimates that ML methods can substantially reduce false declines and improve credit card approvals.

Other ML applications use data sources other than stock prices and trades. In sentiment analysis, ML learning attempts to discover new trends and signals in financial activity and then replicate and enhance human “intuition.”

Hoberg and Lewis (2017) use text-based sentiment analysis of SEC 10-K filings and find that fraudulent firms produce verbal disclosure that is abnormal relative to strong counterfactual evidence. They find evidence that fraudulent managers over purport good performance and disclose fewer details explaining the sources of the firm's performance. Cecchini et al. (2010) use support vector machines for detecting corporate management fraud using basic financial data.

In the case of money laundering detection, a large number of false positive results are often created. As a result, ML in the anti-money laundering area has proved more challenging. Firstly, there is a lack of large public datasets that can be used to analyse money laundering prediction. Many financial institutions still rely on conventional rules-based systems which emphasise individual transactions and simple transaction patterns which may not be sophisticated enough to detect sophisticated transactions (van Liebergen, 2017). Tang and Yin (2005), Liu et al. (2008), Lv et al. (2008) and Villalobos and Silva (2017) detail ML applications to anti-money laundering in a foreign context. Purda and Skillicorn (2015) and Perols et al. (2017) detect fraud from financial statements using ML techniques.

Zhang and Trubey (2018) apply five ML algorithms to US money laundering data. These supervised ML algorithms include Bayes logistic regression, decision tree, random forest, support vector machines and artificial neural networks. Zhang and Turbey (2017) conclude that artificial neural networks consistently outperform parametric logistic regressions. For high event rate regimes with sampling, support vector machines surpass logistic regressions and random forests delivers comparable performance. Clearly ML algorithms provide future possibilities for money laundering detection.

4.2. Banking chatbots and robo-advisory services

The unit cost of financial intermediation in the US has remained at approximately 2% for 130 years (Philippon, 2015). After the 2008 financial crisis, the cost of financial intermediation has declined only marginally in Europe and the US (Philippon, 2016; Bazot, 2013). As part of this post financial crisis response, robo-advisors and chatbots are emerging across the financial services sector, helping consumers choose investments, banking products and insurance policies. A “bot” is a software application created to automate certain tasks using AI technology (Future Today Institute, 2017). A robo-advisor is an algorithm based digital platform that offers automated financial advice or investment management services.

The term “robo-advisor” was essentially unheard-of a decade ago, but it is now relatively commonplace in the financial landscape. However, the term is misleading and doesn't involve robots at all. Instead, robo-advisors are algorithms built to calibrate a financial portfolio to the user's goals and risk tolerance. Chatbots and robo-advisors powered by natural language processing (NLP) and ML algorithms have become powerful tools with which to provide a personalised, conversational and natural experience to users in different domains.

Chatbots and robo-advisors have gained significant appeal with millennial consumers who do not need a physical advisor to feel comfortable investing, and who are less able to validate the fees paid to human advisors.

There are several ways that AI chatbots can improve the banking industry, including helping users manage their money and savings. For example, Plum is a chatbot that can be accessed through Facebook Messenger and helps a customer save money in small increments. After the initial registration, Plum is connected to the customer's bank account and its AI engine then analyses customer income and spending history and then predicts how much they can afford to save. Small amounts are then deposited to the Plum savings account with periodic reporting.

“JP Morgan’s [AI technology] can review approximately 12,000 documents in a matter of seconds...a human would spend 360,000 hours on the same documents”

Brummer and Yadav, 2019

Banks are also engaging chatbots to improve their self-service interfaces. The Bank of America has launched its AI chatbot Erica²⁹ and it is available through voice or message chat on the bank's mobile app. Erica's AI engine also leverages analytics to assist in managing personal finance. JP Morgan has invested in COiN, which is an AI technology that reviews documents and extracts data in far less time than a human. COiN can review approximately 12,000 documents in a matter of seconds, whereas a human would spend more than 360,000 hours of work on the same documents (Brummer and Yadav, 2019).

Chatbots and conversational interfaces are a rapidly expanding area of venture investment and customer service budget. Such chatbots have had to be built with robust natural language processing engines as well as reams of finance-specific customer interactions. Natural language processing is making it increasingly difficult for bank customers to tell whether they are talking to an AI interface or a human. Japan's three megabanks are using AI and robotics to streamline customer questions.³⁰ For example, the Mizuho Group has a robot that helps answer asset management questions and compiles documents.

Amelia is an AI powered chatbot which assists Allstate Insurance employees. Originally deployed in September 2017, Amelia has helped call centre representatives with more than 3 million customer conversations.³¹ Amelia is trained in 40 insurance related topics and uses DL and natural language processing as well as data analytics to understand the intent of the user's text and offer precise answers.

Lemonade is a B2C platform that provides property and casualty insurance to home owners and renters. It uses ML and chatbots for customers. On average, it takes 90 seconds to get insured and three minutes to get paid for a claim. In China, dialogue robots are an AI technology that are currently being utilised in the peer-to-peer (P2P) industry.

An estimated \$200 billion debt mountain has formed in the Chinese P2P industry and since mid-2017 many lenders have shut down as lending controls have been implemented and additional licenses required.

In 2016, Ziyitong was established and launched an AI platform to help recover an estimated RMB150 billion in delinquent loans.³² The AI platform helps recover delinquent loans for approximately 600 debt collection agencies and over 200 lenders (including the Postal Savings Bank of China and Alibaba). The AI platform is a dialogue robot which utilises information about borrowers and their friends' network. The dialogue robot uses the information to determine the phrasing with the highest likelihood of pressuring the borrower to repay the loan. The dialogue robot will also call the borrower's friends, encouraging repayment of the loan. Ziyitong claims its recovery rate is 41% for large clients and loans that are delinquent up to one week, a rate that is twice that of traditional debt collection methods.

Currently the UK and Germany dominate growth in the European robo-advisory industry.³³ However, the US is still considered a major hub of the robo-advisory sector. There are between 98 and 126 robo-advisory services in Europe compared with approximately 200 in the US. Robo-advisors use ETFs in stocks or bonds for portfolio construction. A low interest rate environment has been a major driver of this growth. As interest from bank deposits have become more negligible demand has grown for investment services because robo-advisors provide an inexpensive, convenient platform.

Robo-advisors have the potential to lower costs and increase the quality and transparency of financial advice for consumers. Rohner and Uhl (2017) describe robo-advisory services in three contexts: (1) access to and rebalancing of passive and rule-based investment strategies, (2) cost-efficient implementation of a diversified asset allocation, and (3) behavioural biases. They find that compared to traditional investment advice, robo-advisors can save costs of up to 4.4% per year.

4.3. Algorithmic trading

Algorithmic trading (AT) has become a dominant force in global financial markets. Also called "Automated Trading Systems," AT's origins date back to the 1970's. Kirilenko and Lo (2013) provide a brief survey of the evolution of the AT field. Chakravorty (2016) defines AT as: "Algorithmic trading is about implementing trading rules into a program and using the program to trade, [and AI trading] can be defined as an approach to machine learning that learns the structure of the data, and then tries to predict what will happen".

Algorithmic trading now involves the use of complex AI systems to make extremely fast trading decisions. Computers generate 50-70% of equity market trades, 60% of futures trades and 50% of Treasuries (Brummer and Yadav, 2019). Aldridge and Krawciw (2017) estimate the share of market AT to be closer to 40%.

The benefits of AT include (1) the ability of trades to be executed at the best possible prices, (2) increased accuracy and a reduced likelihood of mistakes, (3) the ability to automatically and simultaneously check multiple market conditions and (4) human errors caused by psychological or emotional conditions are likely to be reduced.

In relation to the second benefit, the European Space Agency's Mosaic Smart Data algorithms are being now deployed to mitigate "fat finger" trades.³⁴ A fat finger trade is where trader accidentally presses the wrong key. Mosaic Smart Data can be deployed to analyse millions of financial trading data points and have recently been high profile cases at Samsung and Deutsche Bank. The algorithms are also being adopted to detect fraud in the financial services industry.

Algorithmic trading's target clientele is hedge funds, proprietary trading houses, bank proprietary trading desks, corporates, and the next generation market makers. AT includes making certain trading decisions, submitting orders, and managing those orders after submission. Biais, Foucault, and Moinas (2011) and Martinez and Roşu (2011) argue that algorithmic speed should have a positive effect on the informativeness of prices. Hendershott et al. (2013) find that AT improves liquidity and enhances the informational content of quotes. On the other hand, AT may also impose higher adverse selection costs on slower trades.

Algorithmic systems often make thousands or millions of trades per day. The term given to this is HFT. HFT is the most recognisable form of AT and uses high-speed communications and algorithms in financial market transactions.³⁵ HFT has both its supporters and detractors. Since 2013, two-thirds of the top 30 cited papers on HFTs show positive market effects from HFTs (Das, 2017). There are supporting arguments that HFT helps with price discovery and efficiency by trading in the direction of permanent price changes and in the opposite direction of transitory pricing errors (Broggard et al. 2014). These types of trading improve market liquidity. Hendershott and Riordan (2013) find that HFT can provide market stability and Menkveld (2014) finds that HFT reduces trading costs. Hasbrouck and Saar (2013) provide evidence that HFT improves market quality and reduces bid-ask spreads. In fact, HFT is changing the traditional field of market microstructure and will continue to be reinvented through new AI and DL techniques.

Most hedge funds and financial institutions do not openly disclose their AI approaches to trading (for proprietary reasons), but it is believed that ML and DL play an important role in calibrating real-time trading decisions. It also involves neural networks, fuzzy logic and pattern recognition. There are four common AT strategies that are now briefly described. *Signal processing* is a mathematical extension of technical analysis based on the art of filtering to eliminate noise and discern trading patterns.

Secondly, there is a strategy known as *market sentiment*. In this strategy the computer is entirely unaware of market activity until it is fed model market data flows and then the algorithm becomes aware of market agitation and participant activity. The objective of market sentiment is to provide the algorithm the appropriate context to analyse and learn market psychology of supply and demand. Third, there is an AT strategy known as *news reader* which does not react to major political events unless it is taught how to artificially read news headlines. Finally, the AT strategy known as *pattern recognition* enables a machine to learn, adapt and react when patterns arise creating revenue opportunities.

US based Sentient Technologies is an AI company that operates a hedge fund and has developed an algorithm that processes millions of data points to find trading patterns and forecast trends. Based on trillions of simulated trading scenarios, Sentient's algorithms use those scenarios to identify and blend successful trading patterns and devise new strategies.

These techniques enable Sentient Technologies to reduce 1,800 days of trading to just a few minutes. Sentient calls its successful trading strategies “genes,” which are then tested in live trading. Here the genes evolve autonomously as they gain experience. Kensho’s algorithm processes millions of market data points to discover correlations and investment opportunities.³⁶

Another hedge fund, Numerai also uses AT to make trading decisions. Instead of developing the algorithms internally, Numerai outsources algorithm development to thousands of anonymous data scientists, who compete to create the best algorithms. In turn, they win cryptocurrency for their efforts. The way in which Numerai shares trading data with the scientists prevents scientists from replicating the fund’s trades while allowing them to build models for better trades. In December 2017, ING launched an AI bond trading tool, Katana.³⁷ Katana shows human traders at the emerging market desk how to gather bond prices more rapidly. After a six-month trial, Katana led to cutting costs by 25% and faster pricing decisions 90% of the time. UBS and JP Morgan have already introduced AI into their trading tools. JP Morgan uses AI algorithms to execute equity trades and UBS has used AI techniques to trade volatility³⁸. AI algorithms are also being used to guide venture capital investments.³⁹.

The practice still has many sceptics, especially with traditional traders who are dubious about the lack of transparency and “black box⁴⁰” nature of AI algorithms. Yadav (2015) states that algorithms used in markets possess “model risk” or the risk that algorithmic programming may not accurately represent the world as it is. Pasquale (2015) states that one drawback is that high frequency traders mimic what other traders are doing without exploring the underlying value of the company whose shares are being traded. AT can also potentially freeze markets and create instability when trading strategies interact in unforeseen ways.

HFT for example has been blamed for the 6 May 2010 “Flash Crash”⁴¹ and is described in Kirilenko et al. (2017), Kirilenko and Lo (2013) and Lewis (2014). On 6 May 2010, the S&P500 plummeted more than 8%, before quickly rebounding. It was described as the first market crash in a new era of automated, algorithmic trading. During the 2016 Brexit referendum, there was a flash crash of the pound sterling by approximately 6% in two minutes.⁴² In 2015, algorithms based on volatility quantitative strategies were attributed to the August market turmoil.⁴³ After market opening on 24 August, the S&P500 triggered circuit breakers nearly 1300 times.

Most empirical research indicates that evidence-based algorithms more accurately predict the future than do human forecasters. What about “algorithm aversion?” This is when choosing between a human forecaster or a statistical algorithm, a person will often choose the human forecaster. Dietvorst et al. (2016, 2015) document that people are especially averse to algorithmic forecasters after seeing them perform, even when they see them outperform a human forecaster. This is because people more quickly lose confidence in algorithmic forecasters compared to human ones after seeing them make the same mistake.

Through an experiment Dietvorst et al. (2016) determine that participants’ preference for modifiable algorithms is indicative of a desire for some control over the forecasting outcome, and not for a desire for greater control over the forecasting outcome. So, participants would prefer an algorithm they could modify.

O'Neill (2017) addresses the potential biases in AT. Kirilenko and Lo (2013) recognise the interconnectedness of global financial markets and the challenges and unintended consequences that brings with it. They list the following unintended consequences: flash crashes, fire sales, careless IPOs, cyber security breaches, and catastrophic trading errors. They call for a more adaptive and systematic approach to regulating AT, where the technological advances of the industry are addressed whilst protecting those who are not as technologically advanced. A UK House of Lords report in April 2018⁴⁴ suggests that the AI sector's full potential would only be realised if potential risks such as algorithmic bias and the opaqueness of "black box" systems can be mollified.

4.4. Other applications of AI

Today, ML has come to play an integral role in many aspects of the financial ecosystem, from approving loans, to managing assets, to assessing risks. The global real estate industry is worth over \$200 trillion (Barnes, 2016) and has been significantly transformed by these new technologies. Proptech is the term that refers to the new and emerging technologies that are disrupting real estate markets. AI is being used as part of the new real estate business models.

Founded in Germany, Leverton is an AI powered data extraction platform that uses DL algorithms to automatically extract key information out of documents such as rental leases, break options and overall clauses. Structured data can be easily accessed from a platform that is understandable in 20 languages.⁴⁵ The platform also has a traceable audit between the underlying documentation and structured data output. In Singapore proptechs use AI to derive formulae calculating the value of a property using a mix of comparative market analysis and algorithms, instead of through manual valuation.⁴⁶

AI algorithms can also be applied to corporate governance contexts. Erel et al. (2018) use ML algorithms to select company directors based on performance. The ML algorithm "learns" from the way directors were selected in the past. They find ML has the potential to improve corporate governance practices because the directors predicted to perform poorly by algorithms, perform much worse than those predicted to perform well. Routledge et al. (2018) investigate phrases associated with mergers and acquisitions (M&A) deal incidence and recent performance. Li et al. (2018) use unsupervised ML to capture corporate culture and its role in M&A activity.

The loan and insurance underwriting business is an ideal environment for ML in finance. At large banks and publicly traded insurance firms ML algorithms can be trained on millions of consumer data items (age, job, marital status, etc.) and financial lending or insurance results (has the person defaulted, paid back the loan on time, been in a car accident, etc.). Using mobile phone data, Bjorkegren and Grissen (2015) employ ML methods to predict loan repayments.

The underlying trends that can be assessed with algorithms, and continuously analysed to detect trends that might influence lending and insuring into the future (are more and more young people in a certain state getting in car accidents)? In the insurance industry, Cytora is using AI to make better risk assessments about their customers, leading to more accurate pricing and minimising claims.⁴⁷

“Machine learning promises to shake up large swathes of finance.”

The Economist, 2017

5. Econometrics versus ML

The goal of both statistics and ML methods is to learn from data (Frame et al., 2018). However, ML methods are not guided by economic theory and are more about algorithms, rather than about asymptotic statistical processes. Traditional statistics highlights hypothesis testing and inference, whereas ML methods emphasise obtaining the best prediction. Unlike maximum likelihood estimation, ML's framework tends to be less unified. To-date there have been relatively limited data intensive applications of ML and DL in the finance literature. However, the increase in processing power, the emergence of big data, better algorithms and growth in Fintech after the 2008 financial crisis have led to an increase in ML and DL techniques and applications.

Most fields (including finance) have traditionally employed models like linear regression where the curve fit to the data is usually a straight line (Domingos, 2017). However, most data tends to exhibit nonlinearity. Several ML methods are able to infer non-linear relationships. The key difference between ML and conventional econometric analysis is its larger focus on prediction compared to summarisation and causal inference (Varian, 2014). ML emphasises “high dimensional prediction problem” and traditional statistics emphasises “formal statistical inference (confidence intervals, hypothesis tests, optimal estimators) in low dimensional problems” (Wall, 2017). Because of this, ML models are not evaluated on the basis of statistical tests, but on their out-of-sample prediction performance.

This means that a ML model describes situations it has not seen before. In order to create a powerful out-of-sample prediction, the model is created by running variables and the model on data subsamples to identify the most powerful predictors. Then testing is conducted (thousands of times) on different data subsamples. This is done so that the model can learn from the data and improve its predictive performance (van Liebergen, 2017). As an example, Khandani et al. (2010) apply ML methods to consumer credit risk models and are able to construct out-of-sample forecasts that significantly improve the classification rates of credit-card-holder delinquencies and defaults. They estimate cost savings ranging from 6% to 25% of total losses.

There are some drawbacks to ML and DL methods. One drawback of ML is that one may struggle to explain why a model is doing what it does, commonly known as the “black box” criticism. The process by which DL techniques reach decisions is also unclear. Deep Learning techniques provide predictions, but they do not provide insight into how the variables are being used to reach these predictions (Wall, 2017). This is especially important for trying to prevent discrimination in lending models (which has been studied in the peer-to-peer lending literature).

Overviews of many ML techniques are provided by Varian (2014), Einav and Levin (2013), Mullainathan and Spiess (2017), Chernozhukov, Chetverikov, Demirer, Duflo, Hansen and Newey (2017), and Chernozhukov, Chetverikov, Demirer, Duflo, Hansen, Newey and Robins (2017). Athey (2015) details how ML relates to causal inference. Many of these reviews provide more technical discussions of various ML methods. ML algorithms are categorised as either supervised learning or unsupervised learning.

The goal of supervised learning is to make predictions by learning patterns from labelled data, whereas in unsupervised learning the aim is to model the underlying structure or distribution in the data to learn more about the data. Generally, ML can also be employed in three types of statistical problems: clustering, classification and regression. Clustering is classified as unsupervised ML and regression and classification are considered supervised ML **Figure 3** (next page) illustrates the distinction between unsupervised and supervised learning methods.

5.1. Unsupervised machine learning

Unsupervised ML is defined as a system that learns without needing human assistance and “labelled” data to teach it. The algorithm aims to find structure in the data, usually by grouping observations, or clusters. Clustering algorithms, dimension reduction, feature extraction and topic models (including natural language processing and text mining models) are examples of unsupervised ML methods. Descriptive techniques such as clustering are often used to gain insights into the mechanisms of certain parts of the financial markets. Unsupervised learning can potentially be quite useful in applications involving text, images, or other very high-dimensional data (Athey and Imbens, 2017).

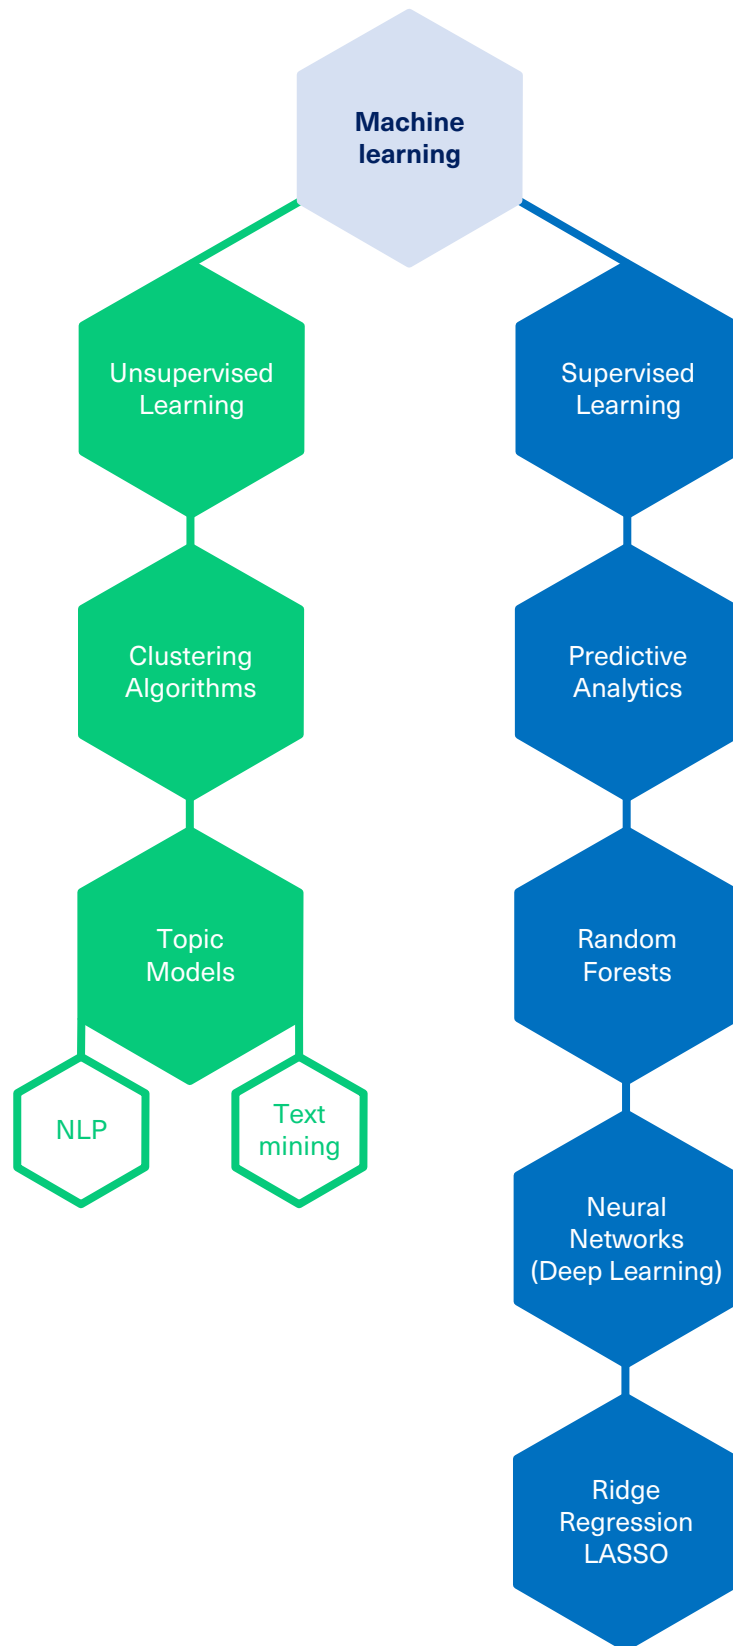
5.1.1. Clustering algorithms

Clustering algorithms are used to understand consumer and trader behaviour. One example is bank overdraft charges in order for bankers to get a better understanding of what is happening to these individuals and what might be the causes of the situation. Regulators can use clustering algorithms to understand traders better (or even insider trading) and categorise business models of firms in advance of regulators' visits.

5.1.2. Topic models

Topic models help us understand the behavioural drivers of different market participants. The area can be further subdivided into natural language processing and textual analysis. One of the earliest applications of ML at the US SEC was to use text mining and natural language processing to detect accounting fraud (Baugess, 2017). Tetlock (2007) uses a popular daily *Wall Street Journal* column to examine the interactions between the media and stock prices. He finds that high media pessimism can predict downward pressure on market prices which followed by a reversion to fundamentals. He finds unusually high or low pessimism predicts high market trading volume.

Figure 3 – Machine learning taxonomy



5.1.2. Topic models (cont.)

Tetlock et al. (2008) use a dictionary-based approach and examine whether language measures can be used to predict individual firms' accounting earnings and stock returns. They find that media content captures otherwise hard-to-quantify aspects of firms' fundamentals, and that investors quickly incorporate this into stock prices. Using ML methods applied to internet messages about stocks, Sinha (2010) argues that the market underreacts to the tone of news articles.

Li (2008) classifies the tone of forward-looking statements, showing its association with firm characteristics using a naïve Bayes algorithm. Manela and Moreira (2017) construct a text-based measure of uncertainty starting in 1890 with WSJ front-page articles. They measure news implied volatility (NVIX) peaks during stock market crashes, times of policy-related uncertainty, world wars, and financial crises. Between 1890 and 2009 a high NVIX predicts high future returns in normal times and rises just before transitions into economic disasters.

Guo et al (2016), Kearney and Liu (2014), Loughran and McDonald (2016) provide a surveys of textual analysis applications in the accounting and finance literature. They observe the challenges of tonal classification to be the use of slang, sarcasm, emoticons and the constantly changing vocabulary on social media. Bao and Datta (2014) use advanced textual analysis to investigate topics in Form 10-Ks related to risks (10-Ks are annual reports required by the SEC). Huang et al. (2018) use this approach to quantify the content of analyst reports and the role of information discovery and interpretation.

Antweiler and Frank (2004) examine whether internet stock message boards can move markets by studying Yahoo! Finance and Raging Bull messages about the companies in the DJIA and the Dow Jones Internet Index. They find that stock messages help predicts market volatility. However, whilst the effect on stock returns is statistically significant, the economic significance is small. Renault (2017) examines the microblogging platform StockTwits and provides empirical evidence that online investor sentiment helps forecast intraday stock index returns.

5.2. Supervised machine learning models

Supervised learning is the case where the chosen algorithm tries to fit the target using the given input features. A set of “training data” that contains labels on the observations is supplied to the algorithm. For example, the labels could be solvent/insolvent. Based on a massive set of data, the algorithm will “learn” a classification rule that it will use to predict the labels for the out of sample observations. For supervised learning methods the test sample might have 10% of observations (Athey and Imbens, 2017).

The difference between supervised and unsupervised ML is that the latter lacks labelled training data and has to determine correlations by itself. Examples of supervised learning include predictive analytics, random forests, neural networks and LASSO techniques. Introduced by Tibshirani (1996), LASSO stands for least absolute shrinkage and selection operator. Belloni et al. (2014) provide a very thorough discussion of how LASSO techniques can be applied to high dimensional data.

Reinforcement learning has its roots in the pioneering work of Bellman (1957) on optimal control and falls between unsupervised and supervised machine learning (FSB, 2017). In reinforcement learning the connection between actions and rewards is unknown and must be inferred from the agent's interactions with the environment. An unlabelled set of data is fed to an algorithm and then the algorithm chooses an action for each data point. A human may then provide feedback (the supervised part) that subsequently helps the algorithm learn. Corazza and Bertoluzzo (2014) and Ritter (2017) apply reinforcement learning to dynamic trading strategies.

5.2.1. Predictive analytics

Both clustering algorithms and topic models are precursors to predictive analytics, a form of supervised machine learning. This is where one teaches an algorithm to learn from past breaches of regulations and to detect new breaches in regulation. This has potential applications in cartel detection and insider trading.

5.2.2. Random forests

The random forest model was developed by Breiman (2001). The technique is based on decision tree models and is also known as generalised classification and regression trees (CART). Breiman (2001), Booth et al. (2014) and Bagherpour (2018) provide a detailed technical overview of random forest methodology. The random forest model applies the bootstrap aggregating technique to tree learners. The random forest algorithm is extremely useful for forecasting as well as allowing classification on whether a variable is more or less important when constructing the decision tree.

Yeh et al. (2014) find that random forest techniques are very robust and allow for the presence of outliers and noise in the training set. JPMorgan researchers consider that random forest shows promise for trading 10-year US Treasury market instruments. Medeiros et al. (2017) recognise the nonlinearity property with inflation dynamics and compare 16 ML methods against benchmark statistical models. In general, ML methods with a large set of covariates provide superior results compared to univariate benchmarks and factor models. Medeiros et al. (2017) find that random forests is the best model indicating a degree of nonlinearity in the dynamics of inflation. Using ML applications in bankruptcy prediction, Barboza et al. (2017) find that random forest techniques outperform other methods.

5.2.3. Neural networks

In 2012 the full potential of neural networks (NN) applications in finance first came to wider attention. NN are a form of DL methodology. Das (2017) observes that DL has been proven to uncover subtle nonlinearities in data that are not discoverable using standard linear econometric models. General NN consist of three types of layers: input layer, a hidden layer and output layer, and each layer also consist of multiple units. Lacher et al. (1995) use NN to assess corporate financial health. They find that NN overcome the limitations of multi-discriminant analysis (MDA) such as multivariate normality, linear reparability and independence of the predictive variables and achieve superior results.

Altman et al. (1994) compare artificial neural networks (ANN) with traditional statistical methodologies including logit analysis and discriminant analysis. They analyse the financial health of 1,000 Italian industrial firms between 1982 and 1992 and find NN display acceptable classification and holdout sample accuracy. Using S&P data between 1970 and 1989, Coats and Fant (1999) find that NN are more effective than MDA for early detection of financial distress and identify firms which receive opinions regarding whether the firm is still a going concern.

Other papers that show NN techniques are at least as good as MDA methods include: Williams (1985); Odom and Sharda (1990); Utans and Moody (1991) when applied to bond ratings predictions. NN performs well with regard to limited availability of data and by the lack of complete a priori information that could be used to impose a structure to the network architecture. Garavaglia (1991) also looks at bond ratings using NN. Salchenberger et al. (1992) compare NN with MDA techniques using thrift failure applications. They find that NN require fewer assumptions and achieve a higher degree of prediction, accuracy and is more robust than traditional statistical methods.

Tama and King (1992) compare NN methods with linear classifiers and logistic regression by applying these to Texan bankruptcies. Angelini et al. (2008) apply NN to model credit risk for Italian small and medium sized enterprises and find that NN methods perform very well. Hutchinson et al. (1994) use NN algorithms and find it is an excellent framework to model options pricing functions. Amilon (2003) employs daily call option pricing data to examine whether NN can outperform the Black-Scholes formula. NN are applied to out of sample pricing and delta hedging and Amilon (2003) finds NN outperform the Black-Scholes formula both in hedging and pricing performance.

Culkin and Das (2017) employ a fully-connected DL NN to reproduce the Black and Scholes (1973) option pricing formula. First, Culkin and Das (2017) start with the application of learning the Black-Scholes option pricing model from simulated data. The next step is that market data is used to train an option pricing model that performs better in terms of predictive accuracy, adaptability and robustness, which is their main result. This is achieved with a high degree of accuracy. Li et al. (2018) employ DL—specifically, a method that uses layers of ANN to automatically process earnings calls and score cultural values for companies. Their ML approach yields a high-quality culture dictionary useful for measuring corporate culture which is also scalable to a large collection of textual data.

5.2.4. Support vector machine (SVM)

Support vector machines are related to and contain elements of ML, NN and non-parametric applied statistics. Bagherpour (2018) provides an excellent detailed technical discussion of the SVM methodology. Cecchini et al. (2010) use SVMs for detecting corporate management fraud using financial data. Société Générale, the French bank, on the other hand, says support-vector machines are better, over time, for making long/short equity investment decisions. Hardle et al. (2008) employ smooth SVMs to predict firm default risk. Hardle, Moro and Hoffmann (2011) find SVMs are capable of extracting the necessary information from financial balance sheets and then to predict the future company solvency or insolvency.

6. Machine learning versus quantum computing

Quantum computing (QC) is a relatively new field of research that studies the algorithms and systems that apply quantum phenomena to complex problems. QC can potentially process data at speeds that are impossible for traditional computers. Traditional computers only process information in a binary format: in zeroes or ones. Benioff and Feynman established the field in the early 1980s and noted that whereas digital computers (DC) could not efficiently simulate a probabilistic system, but QC could.

QCs are much better suited than DCs to solve financial problems because QCs operate with random variables, whereas DCs just simulate random variables. QC can hold multiple states simultaneously and these co-existing states are called qubits. Qubits are able to process four values at any given time and allows the computer to parallel process information (Future Today Institute, 2017).

The fact that qubits are memory elements may hold a linear superposition of both states is a game changer for the financial services industry. Rolling IT security is one of the most famous applications of QC. Rosenberg et al. (2015) apply QC to portfolio optimisation problems. Lopez de Prado (2016) suggests that QC has potential with finance problems that involve scenario analysis or option pricing.

In the case of scenario analysis, QC can evaluate an extremely large number of outcomes that have been generated at random and for option pricing QC can evaluate a large number of paths that can be computationally expensive. Looking ahead, ML is expected to have a far more powerful impact if it is combined with QC capabilities. Lopez de Prado (2016) describes how QC has potential to more efficiently refine clustering algorithms.

7. Regulation and policy-making

AI and ML are moving faster than policy makers can understand to the extent it is almost outstripping the current legal and regulatory framework. Technology is opaque and fast moving and regulators find it hard to keep pace, for both the cumulative impact and risks of contagion. Athey and Imbens (2017) and Mullainathan and Spiess (2017) argue that ML methods hold great promise for improving the credibility of policy evaluation.

The technology underpinning Fintech is also fuelling a spinoff field known as RegTech which aims to make compliance and regulatory activities easier, faster and more efficient. RegTech utilises Big Data and ML. RegTech is an emerging field to reduce costs and increase effectiveness. Alarie, Niblett and Yoon (2016) explore how ML technology can improve regulation of human behaviour. They argue that ML techniques can provide fast, accurate and consistent judgements, and streamline operations with reduced error. Baracos and Selbst (2016) examine how fairness and non-discrimination can be implemented in learning programs used in banking, insurance and the justice system.

One example of RegTech can be found at the Hong Kong Stock Exchange which is employing AI software used by NASDAQ to detect stock manipulation and market abuses. The software analyses historical trading activity to identify patterns and anomalies, including dramatic swings in stock prices and rises in trading volume.⁴⁸

Financial regulators are also exploring the use of AI for better monitoring of financial institutions. The UK Financial Conduct Authority (FCA) is examining the possibility of making its handbook machine-readable and then fully machine-executable. This would mean that machines can interpret and implement the rules directly (Citi, 2018). The Division of Economic and Risk Analysis (DERA) at the SEC is exploring ML to extract actionable insights from massive datasets, helping examiners find cases of potential fraud or misconduct (Baugess, 2017). The Monetary Authority of Singapore is examining how AI/ML methods can be used to identify suspicious transactions. This will allow Singapore regulators to better devote resources to higher risk transactions (FSB Report, 2017).

International regulators utilise AI-supported analytical methods to recognise vulnerability patterns, scan lengthy reports or analyse incoming data. In 2017 the Bank of England (BoE) joined forces with MindBridge to use an AI auditor to help detect anomalies in transactions and reports. The BoE will be utilising ML to achieve these goals.⁴⁹

Chakraborty and Joseph (2017) outline how ML methodologies can be applied in the context of central banking. They provide three ML case studies (two supervised learning cases and one unsupervised learning case) to compare ML methods against traditional econometric methods. They examine ANN, tree-based models, support vector machines, recommender systems and different clustering techniques.

Chakraborty and Joseph (2017) conclude that ML methods generally outperform traditional modelling approaches in prediction tasks. However, as regards causal inference properties, open research questions remain.

In 2018, Chancellor Angela Merkel announced that the German government would spend €3 billion to boost AI capabilities. The Deutsche Bundesbank is already using AI in its risk management area and uses NN to assess financial market soundness. On 25 May 2018 the General Data Protection Regulation (GDPR) came into force. Under GDPR, EU citizens have the right to receive an explanation for decisions based solely on automatic processing. Under recital 71 of the GDPR, “the data subject should have the right not to be subject to a decision, which may include a measure, evaluating personal aspects relating to him or her which is based solely on automated processing and which produces legal effects concerning him or her or similarly significantly affects him or her, such as automatic refusal of an online credit application or e-recruiting practices without any human intervention.”

Furthermore, GDPR stipulates that companies must first obtain consent from an EU citizen before using consumer data. If the EU citizen’s data is stored on servers located outside of the EU region, GDPR rules apply. Failure to comply to GDPR can result in substantial fines: up to \$22 million or 4% of a company’s revenues (Deloitte, 2018).

The European MIFID II⁵⁰ (which also came into effect in 2018) requires that firms applying algorithmic models based on AI and ML should have a robust development plan in place. Firms need to ensure that potential risks are included at every stage of the process (Wuermaling, 2018). In February 2018 the FCA and Prudential Regulatory Authority released consultation papers on algorithmic trading which lists key areas of supervisory focus in relation to MIFID II.

Prior to the 2008 financial crisis, ratings agencies used algorithmic assessments of creditworthiness to rate mortgage backed securities (MBS). The 2010 Dodd Frank Act requires ratings agencies to disclose material changes in ratings agencies methodologies. The US SEC has been testing ML applications in regulatory settings (Bauguess, 2017). Early use of AI at the SEC were simple text analytic methods used to determine if CDS risks in the run up to the global financial crisis could have been anticipated (Bauguess, 2017). The first mention of a CDS contract in a US 10K filing was in 1998. By 2004, more than 100 corporate issuers mentioned their use in 10Ks. A large increase in disclosures occurred in 2009 after the crisis was in full swing. Even though this early attempt did not work as planned, it was an early example of proof of concept and structuring certain phrases.

The next stage the SEC pursued was to use natural language processing efforts to measure the probability of certain words in SEC documents (or unsupervised learning). Topic models were deployed to analyse the information in the tips, complaints, and referrals received by the SEC. Additionally, tonality analysis is employed to gauge the negativity of a filing by counting the appearance of certain financial terms that have negative connotations. The objective is to learn whether themes can be classified directly from the data itself that would enable more efficient allocation of resources.

In 2013 the SEC disclosed that they were using a natural language processing model to detect accounting fraud. The SEC now employs a Corporate Issuer Risk Assessment Program which provides 230 metrics to staff in identifying fraud and misconduct (Hunt, 2017). Now for many cases the SEC utilises a two-stage process: unsupervised learning methods to identify general themes and then supervised ML algorithms to predict the presence of idiosyncratic risks. Bauguess (2017) observes that back testing of the ML methods performs very well.

The UK Serious Fraud Office (SFO) includes AI in its work. In a typical year the SFO processes over 100 million documents in fraud and corruption cases. This includes the Rolls-Royce bribery case, which resulted in the largest ever fine imposed in the UK for criminal conduct.⁵¹ The SFO used the RAVN robotic system, which ended up costing £50,000, and saved UK taxpayers hundreds of thousands of pounds. RAVN is referred to as a Legal Professional Privilege (LPP) robot. RAVN sifts documents into “privileged” versus “non-privileged” piles, indexes and compiles summaries. In the Rolls-Royce case, RAVN processed 30 million documents at a rate of up to 600,000 per day (compared with a team of lawyers that would have processed 3,000 per day). The SFO is also considering other advanced ML techniques in other aspects of its operations.

Wuermaling (2018) provides six warnings for central bankers in regard to AI: (1) Don't miss out on opportunities of AI in finance; (2) beware of the risks (3) consumers should take care: they remain risk takers (4) Fintechs should not ignore the legitimate concern of society and supervisors (5) AI needs new forms of supervision and (6) central banks should embrace AI.

“The rise of powerful AI will be either the best or the worst thing ever to happen to humanity. We do not yet know which.”

Stephen Hawking

8. Conclusion and directions for future research

Artificial intelligence in the financial services industry is still in early days. AI will become more ubiquitous in finance, and with that comes more challenges including legal, ethical, economic and social hurdles. AI will also continue to bring new complexities to the global financial ecosystem. As more and more data become available and computing power increases, AI programs will become more complex. But are AI, ML and DL so different from previous advancements they could upend the laws of finance? ⁵²

A fascinating 2018 *Financial Times* article raises this question. For example, in 2017, depreciation represented approximately one-third of global operating cash flows. But what happens to a computer with AI that actually becomes smarter over time? Machine learning by definition is the opposite of decay.⁵³ So if an asset grows more valuable with use then it should be negatively depreciated. This means for certain technology firms the earnings boost would potentially be quite large. Artificial intelligence combined with the internet of things (IoT) will result in physical things becoming more adaptive and responsive which will extend their useful lives.

Along with big data, AI is viewed in the financial services sector as a technique that has the potential to deliver huge analytical power. Yet many risks still need to be addressed. Many AI techniques remain untested in financial crisis scenarios. There have been several instances in which the algorithms implemented by financial firms appeared to act in ways quite unforeseen by their developers, leading to errors and flash crashes (notably the pound's flash crash following the Brexit referendum in 2016). Lo (2016) calls for developing more robust technology capable of adapting to human foibles so that users can employ these tools safely, effectively and effortlessly.

Much remains to be done. And there is clearly a need for more education on AI literacy and awareness⁵⁴. The late Stephen Hawking summed it up: “The rise of powerful AI will be either the best or the worst thing ever to happen to humanity. We do not yet know which.”

Appendix

Timeline of artificial intelligence milestones

- 1937** Claude Shannon proposes that Boolean algebra can be used to model electronic circuits
- 1943** McCulloch & Pitts recognise that Boolean circuits can be used to model brain signals
- 1950** Alan Turing develops the Turing Test
- 1950** Minsky and Edmonds build the first neural network computer (the SNARC)
- 1956** The term "*artificial intelligence*" is coined by John McCarthy
- 1956** Newell and Simon create the Logic Machine
- 1957** Economist Herbert Simon predicts that computers would defeat humans at chess within the following decade
- 1958** Frank Rosenblatt introduces a new form of neural network known as "perceptron"
- 1958** Early genetic algorithms experiments
- 1959** Arthur Samuels demonstrates that a computer can play checkers better than its creator, and even play against itself to practice
- 1961** Newell and Simons creates General Problem Solver
- 1964** Computers understand natural language enough to solve algebraic and word problems
- 1965** Herbert Dreyfus' report severely criticises the emerging AI field
- 1967** Marvin Minsky predicts that within a generation the problem of creating "artificial intelligence" would be solved
- 1969** Bryson and Ho develop a back propagation algorithm
- 1971** Terry Winograd's program SHRDLU answers questions in natural language
- 1973** UK Lighthill Report ends British government support for AI research
- 1974 – 1980** **First "AI Winter"**
- 1980** Expert Systems, or Knowledge Systems, emerge as a new field within AI
- 1980s** Early part of decade – Benioff and Feynman create Quantum Computing
- 1982** PlanPower is conceptualised by Applied Expert Systems (APEX)
- 1982** James Simons starts quant investment firm Renaissance Technologies
- 1984** American Association for AI coins the term "AI Winter"
- 1987** Personal Financial Planning System (PFPS) used by Chase Lincoln First Bank

1987 – 1993 Second “AI Winter”

- 1988** David Shaw founds D.E. Shaw and is an early adopter of AI among its hedge funds
- 1990s** The AI industry shows renewed interest in neural networks
- 1990** Neural net device reads handwritten digits to determine amounts on bank cheques
- 1993** FinCen puts FAIS (its AI system) into service to monitor money laundering
- 1997** Deep Blue defeats Garry Kasparov, world chess champion at the time. IBM's stock price increases by \$18 billion
- 2005** The DARPA 132-mile challenge sees AI applied to autonomous driving
- 2007** The DARPA Urban Challenge
- 2009** Google's first self-driving car
- 2010** Flash Crash occurs on 6 May. In 36 minutes, the S&P crashed 8%, before a rebound
- 2012** On 1 August, Knight Capital loses \$440 million 45 minutes after deploying unverified trading software
- 2014** Man Group starts to use AI to manage client money
- 2016** Google's DeepMind AlphaGo applies ML algorithms to win at international Go championship
- 2017** Two Sigma hedge fund which uses ML, crosses the \$50 billion in assets under management
- 2017** Beijing announces plans to lead the world in AI by 2030
- 2018** UBS announces development of recommendation algorithms
- 2018** The Merkel government announces €3 billion will be spent on AI capabilities
- 2018** President Macron announces that all algorithms developed for government use will be made publicly available
- 2018** Alibaba announces plans to bring AI chips to market the following year
- 2018** MiFID II takes effect
- 2018** GDPR takes effect on 25 May
- 2018** Baidu becomes the first Chinese tech giant to join a US led consortium on AI safeguards

References

- Alarie, Benjamin and Niblett, Anthony and Yoon, A. (2016) Regulation by Machine. Available at SSRN: <https://ssrn.com/abstract=2878950>
- Altman, E. I., Marco, G., & Varetto, F. (1994). Corporate distress diagnosis: Comparisons using linear discriminant analysis and neural networks (the Italian experience). *Journal of Banking & Finance*, 18(3), 505-529.
- Amilon, H. (2003). A neural network versus Black–Scholes: a comparison of pricing and hedging performances. *Journal of Forecasting*, 22(4), 317-335.
- Angelini, E., di Tollo, G., & Roli, A. (2008). A neural network approach for credit risk evaluation. *The quarterly review of economics and finance*, 48(4), 733-755.
- Antweiler, W., & Frank, M. Z. (2004). Is all that talk just noise? The information content of internet stock message boards. *The Journal of finance*, 59(3), 1259-1294.
- Athey, S. (2017). The Impact of Machine Learning on Economics. In *Economics of Artificial Intelligence*. University of Chicago Press.
- Athey, S. (2015). Machine learning and causal inference for policy evaluation. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 5-6). ACM.
- Athey, S., & Imbens, G. W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, 31(2), 3-32.
- Auria, Laura and Moro, R. A. (2008) Support Vector Machines (SVM) as a Technique for Solvency Analysis. DIW Berlin Discussion Paper No. 811. Available at SSRN: <https://ssrn.com/abstract=1424949> or <http://dx.doi.org/10.2139/ssrn.1424949>
- Bagherpour, A. (2018) Predicting Mortgage Loan Default with Machine Learning Methods. Working Paper.
- Baker, Tom and Dellaert, Benedict G. C., Regulating Robo Advice Across the Financial Services Industry (2018). *Iowa Law Review*, Vol. 103, P. 713, 2018; U of Penn, Inst for Law & Econ Research Paper No. 17-11.
- Bank of England (2018) Algorithmic Trading. Bank of England Consultation Paper.
- Bao, Y., and A. Datta, (2014) Simultaneously discovering and quantifying risk types from textual risk disclosures, *Management Science* 60, 1371–1391.
- Baracos, S. & Selbst, A. (2016) Big Data's Disparate Impact. *California Law Review*, 104(3):671-732
- Barboza, F., Kimura, H., & Altman, E. (2017). Machine learning models and bankruptcy prediction. *Expert Systems with Applications*, 83, 405-417.

- Barnes, Y. (2016) Around the world in dollars and cents. Available at: http://www.savills.co.uk/research_articles/188297/198667-0.
- Bauguess, S. W (2017) The Role of Big Data, Machine Learning, and AI in Assessing Risks: A Regulatory Perspective, Keynote Address: OpRisk North America.
- Bazot, G. (2013) Financial consumption and the cost of finance: Measuring financial efficiency in Europe (1950-2007). Working Paper. Paris School of Economics.
- Behan, J., & Lecot, K. (1987) Overview of financial applications of expert systems. In *Proceedings of WESTEX-87: IEEE Western Conference on Knowledge-Based Engineering and Expert Systems* (pp. 223-229).
- Bellman, R. (2013) *Dynamic programming*. Courier Corporation.
- Biais, B., Foucault, T., & Moinas, S. (2011). Equilibrium high frequency trading. In *International Conference of the French Finance Association (AFFI)*.
- Björkegren, Daniel and Grissen, Darrell, (2018) Behavior Revealed in Mobile Phone Usage Predicts Loan Repayment. Available at SSRN: <https://ssrn.com/abstract=2611775> or <http://dx.doi.org/10.2139/ssrn.2611775>
- Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637{659.
- Booth, A., Gerding, E., & McGroarty, F. (2014). Automated trading with performance weighted random forests and seasonality. *Expert Systems with Applications*, 41(8), 3651-3661.
- Brandi, J. T. (1988). Artificial Intelligence Expert Systems and Financial Planning. *Journal of Financial Service Professionals*, 42(2), 64.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1), 5-32.
- Brogaard, J., Hendershott, T., & Riordan, R. (2014). High-frequency trading and price discovery. *The Review of Financial Studies*, 27(8), 2267-2306.
- Brown, C. E., Nielson, N. I., & Phillips, M. E. (1990). Expert systems for personal financial planning. *Journal of Financial Planning*, 3(3).
- Brummer, C., & Yadav, Y. (2019). The Fintech Trilemma. *Georgetown Law Journal*, 107, 235-307.
- Buchanan, B and C. Cao (2018) Quo Vadis? Fintech in China Versus the West. Working Paper. Available at: https://swiftinstitute.org/wp-content/uploads/2018/10/SIWP-2017-002-_Fntech_China_West_BuchCao_FINAL.pdf.
- Burges, C. J. (1998). A tutorial on support vector machines for pattern recognition. *Data mining and knowledge discovery*, 2(2), 121-167.

- Butaru, F., Chen, Q., Clark, B., Das, S., Lo, A. W., & Siddique, A. (2016). Risk and risk management in the credit card industry. *Journal of Banking & Finance*, 72, 218-239.
- CB Insights (2018) Top AI Trends to Watch in 2018. Available at: <https://www.cbinsights.com/research/report/artificial-intelligence-trends-2018/>
- Cecchni, M., Aytug, H., Koehler, G. J. and Pathak, P. (2010) Detecting management fraud in public companies. *Management Science*, 56 (7), 1146-1160
- Chakraborty, C., & Joseph, A. (2017). Machine learning at central banks. Available at: <https://www.bankofengland.co.uk/working-paper/2017/machine-learning-at-central-banks>.
- Chakravorty, G. (2016) What is the difference between AI and algo trading? Available: <https://www.quora.com/What-is-the-difference-between-AI-trading-and-algo-trading>.
- Chaboud, A. P., Chiquoine, B., Hjalmarsson, E., & Vega, C. (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance*, 69(5), 2045-2084.
- Cheeseman, P. C. (1985). In Defense of Probability. In *IJCAI* (Vol. 2, pp. 1002-1009).
- Chen, K. C., & Liang, T. P. (1989). PROTRADER: An expert system for program trading. *Managerial Finance*, 15(5), 1-6.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., & Newey, W. (2017). Double/Debiased/Neyman machine learning of treatment effects. *American Economic Review*, 107(5), 261-65.
- Citi (2018) Bank of the Future: the ABCs of Digital Disruption in Finance. CitiReport March 2018.
- Coats, P. K., & Fant, L. F. (1993). Recognizing financial distress patterns using a neural network tool. *Financial Management*, 142-155.
- Corazza, M and Bertoluzzo, F. (2014) Q-Learning-Based Financial Trading Systems with Applications. University Ca' Foscari of Venice, Dept. of Economics Working Paper Series No. 15/WP/2014. Available at SSRN: <https://ssrn.com/abstract=2507826> or <http://dx.doi.org/10.2139/ssrn.2507826>
- Culkin, R. and S. Das (2017) Machine Learning in Finance: The Case of Deep Learning in Option Pricing. Working Paper.
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4), 303-314.
- Das, S.R (2017) The Future of FinTech. Available at: <https://srdas.github.io/Papers/fintech.pdf>.
- Deloitte Insights (2018) Tech Trends: The Symphonic Enterprise. Available at: https://www2.deloitte.com/content/dam/insights/us/articles/Tech-Trends-2018/4109_TechTrends-2018_FINAL.pdf.

- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2016). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114.
- Domingos, P. (2015). *The master algorithm: How the quest for the ultimate learning machine will remake our world*. Basic Books.
- Dreyfus, Hubert L (1965). *Alchemy and Artificial Intelligence*. Santa Monica, CA: RAND Corporation. Available at: <https://www.rand.org/pubs/papers/P3244.html>.
- Erel, I., C. Tan, L. H. Stern, and M. S. Weisbach, (2018) *Selecting directors using machine learning*, Ohio State University working paper.
- Financial Stability Board (2017) *Artificial intelligence and machine learning financial services*. Working Paper. Accessed at: <http://www.fsb.org/2017/11/artificial-intelligence-and-machine-learning-in-financial-service/>
- Finnerty, J. E., & Park, H. Y. (1987). Empirical evidence on stock index arbitrage: the case of program trading. *BEBR faculty working paper; no. 1321*.
- Frame, W. S., Wall, L. D., & White, L. J. (2018). Technological Change and Financial Innovation in Banking: Some Implications for Fintech. In *Oxford Handbook of Banking*, edited by Berger, A., Molyneux, P and J.O. Wilson. 3rd Edition.
- Future Today Institute (2017) *Tech Trends annual report*. Available at: <https://futuretodayinstitute.com/2017-tech-trends/>.
- Garavaglia, S. (1991). An application of a counter-propagation neural network: simulating the Standard and Poor's Corporate Bond Rating system. In *Artificial Intelligence Applications on Wall Street, 1991. Proceedings., First International Conference on* (pp. 278-287). IEEE.
- GDPR (2018) https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en
- Gilmartin, K. J., Newell, A., & Simon, H. A. (1976). A program modeling short-term memory under strategy control. *The structure of human memory*, 15-30.
- Golberg et al (1995) *The FinCEN Artificial Intelligence Systems: Identifying Potential Money Laundering from Reports of Large Cash Transactions*
- Golmohammadi, K., & Zaiane, O. R. (2015). Time series contextual anomaly detection for detecting market manipulation in stock market. In *Data Science and Advanced Analytics (DSAA), 2015. 36678 2015. IEEE International Conference on* (pp. 1-10). IEEE.
- Golmohammadi, K., & Zaiane, O. R. (2012). Data mining applications for fraud detection in securities market. In *Intelligence and Security Informatics Conference (EISIC), 2012 European* (pp. 107-114). IEEE.

- Guo, L., Shi, F., & Tu, J. (2016). Textual analysis and machine learning: Crack unstructured data in finance and accounting. *The Journal of Finance and Data Science*, 2(3), 153-170.
- Hall, W and Pesenti, J (2017) Growing the Artificial Intelligence Industry in the UK, Working Paper. Available at:
https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf.
- Hand, D. J., & Henley, W. E. (1997). Statistical classification methods in consumer credit scoring: a review. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 160(3), 523-541.
- Härdle, W. K., Hoffmann, L., & Moro, R. (2011). Learning machines supporting bankruptcy prediction. In *Statistical Tools for Finance and Insurance* (pp. 225-250). Springer, Berlin, Heidelberg.
- Härdle, Wolfgang K. and Lee, Yuh-Jye and Schäfer, Dorothea and Yeh, Yi-Ren (2007). The Default Risk of Firms Examined with Smooth Support Vector Machines. SFB 649 Discussion Paper 2008-005. Available at SSRN: <https://ssrn.com/abstract=2894311>.
- Heaton, J. B., Polson, N. G., & Witte, J. H. (2017). Deep learning for finance: deep portfolios. *Applied Stochastic Models in Business and Industry*, 33(1), 3-12.
- Hendershott, T., & Riordan, R. (2013). Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis*, 48(4), 1001-1024.
- Hall, W and J. Pesenti (2017) Growing the Artificial Intelligence Industry in the UK. Available at:
<https://www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in-the-uk>.
- Hoberg, G., & Lewis, C. (2017). Do fraudulent firms produce abnormal disclosure? *Journal of Corporate Finance*, 43, 58-85.
- Hoberg, G., & Maksimovic, V. (2014). Redefining financial constraints: A text-based analysis. *The Review of Financial Studies*, 28(5), 1312-1352.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press. First edition.
- Huang, A. H., R. Lehavy, A. Y. Zang, and R. Zheng (2018) Analyst information discovery and interpretation roles: A topic modeling approach, *Management Science*. 64, 2473-2472.
- Hunt, S. (2017) From Maps to Apps: The Power of Machine Learning and Artificial Intelligence for Regulators. Beesley Lecture Series on Regulatory Economics.
- Hutchinson, J. M., Lo, A. W., & Poggio, T. (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks. *The Journal of Finance*, 49(3), 851-889.
- Javelin Strategy (2015) Future Proof Card Authorization. Working Paper. Available at:
<https://www.javelinstrategy.com/press-release/false-positive-card-declines-push-consumers-abandon-issuers-and-merchants>.

- Kaplan, J. (2016). *Artificial Intelligence: What Everyone Needs to Know*. Oxford University Press.
- Kearney, C., & Liu, S. (2014). Textual sentiment in finance: A survey of methods and models. *International Review of Financial Analysis*, *33*, 171-185.
- Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, *34*(11), 2767-2787.
- Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The Flash Crash: High-frequency trading in an electronic market. *The Journal of Finance*, *72*(3), 967-998.
- Kirilenko, A. A., & Lo, A. W. (2013). Moore's law versus murphy's law: Algorithmic trading and its discontents. *Journal of Economic Perspectives*, *27*(2), 51-72.
- Koyuncugil, A. S., & Ozgulbas, N. (2012). Financial early warning system model and data mining application for risk detection. *Expert Systems with Applications*, *39*(6), 6238-6253.
- Lacher, R. C., Coats, P. K., Sharma, S. C., & Fant, L. F. (1995). A neural network for classifying the financial health of a firm. *European Journal of Operational Research*, *85*(1), 53-65.
- Lessmann, S., Baesens, B., Seow, H. V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, *247*(1), 124-136.
- Lewis, M., & Baker, D. (2014). *Flash boys*. New York: WW Norton.
- Li, K., Mai, F., Shen, R., & Yan, X. (2018). Measuring Corporate Culture Using Machine Learning. Available at SSRN 3256608.
- Li, F. (2008) Annual report readability, current earnings, and earnings persistence, *Journal of Accounting and Economics* *45*, 221-247.
- Liu, X., Zhang, P., Zeng, D. (2008) Sequence matching for suspicious activity detection in anti-money laundering, Proceedings of the IEEE ISI 2008 PAISI, PACCF, and SOCO international workshops on Intelligence and Security Informatics.
- Lo, Andrew W., Moore's Law vs. Murphy's Law in the Financial System: Who's Winning? (2016). BIS Working Paper No. 564. Available at SSRN: <https://ssrn.com/abstract=2789737>
- Loughran, T., & McDonald, B. (2016). Textual analysis in accounting and finance: A survey. *Journal of Accounting Research*, *54*(4), 1187-1230.
- Loughran, T., & McDonald, B. (2014). Measuring readability in financial disclosures. *The Journal of Finance*, *69*(4), 1643-1671.
- Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, *66*(1), 35-65.

- Lopez de Prado, Marcos and Lewis, Michael J. (2018) Detection of False Investment Strategies Using Unsupervised Learning Methods. Available at SSRN: <https://ssrn.com/abstract=3167017>
- Lopez de Prado, Marcos, Financial Quantum Computing (2016). Available at SSRN: <https://ssrn.com/abstract=2848632> or <http://dx.doi.org/10.2139/ssrn.2848632>
- Lv, L.-T., Ji, N., Zhang, J.-L. (2008) A RBF neural network model for anti-money laundering, *Proceedings of the 2008 International Conference on Wavelet Analysis and Pattern Recognition*.
- Magnuson, William J., Regulating Fintech (2017). *Vanderbilt Law Review*, Forthcoming; Texas A&M University School of Law Legal Studies Research Paper No. 17-55. Available at SSRN: <https://ssrn.com/abstract=3027525>
- Makridakis, S. (2017). The forthcoming artificial intelligence (AI) revolution: Its impact on society and firms. *Futures*, 90, 46-60.
- Manela, A., & Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of Financial Economics*, 123(1), 137-162.
- Martinez, V. H., & Rosu, I. (2013). High frequency traders, news and volatility. In *AFA 2013 San Diego Meetings Paper*.
- McCulloch, W.S. and W. Pitts (1943) A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115-137.
- Medeiros, M., Vasconcelos, G., Veiga, A. and Zilberman, E. (2017) Forecasting Inflation in a Data Rich Environment: The Benefits of Machine Learning Methods. Big Data, Machine Learning and the Macro economy Conference. Norges Bank.
- Mizuta, T., Izumi, K., & Yoshimura, S. (2013). Price variation limits and financial market bubbles: Artificial market simulations with agents' learning process. In *Computational Intelligence for Financial Engineering & Economics (CIFER), 2013 IEEE Conference*.
- Mullainathan, S., & Spiess, J. (2017). Machine learning: an applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87-106.
- Newell, A., Shaw, J.C. and H.A. Simons (1957) Empirical explorations with the logic theory machine. Proc. Western Joint Computer Conference, 15, 218-239.
- Newell, A. and H.A. Simons (1961) GPS, a program that stimulates human thought. In Billing. H. (Ed.) *Lernenende Automaten*, pp. 109-124.
- Odom, M. D., & Sharda, R. (1990). A neural network model for bankruptcy prediction. In *Neural Networks, 1990., 1990 IJCNN International Joint Conference on* (pp. 163-168). IEEE.
- O'Neil, C. (2017). Weapons of math destruction: How big data increases inequality and threatens democracy. Broadway Books.

- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- Pearl, J. (1988). Morgan Kaufmann series in representation and reasoning. Probabilistic reasoning in intelligent systems: Networks of plausible inference.
- Perols, J. L., R. M. Bowen, C. Zimmermann, and B. Samba (2017). Finding needles in a haystack: Using data analytics to improve fraud prediction, *Accounting Review* 92, 221-245.
- Philippon, T. (2016) The Fintech Opportunity. NBER Working Paper 22476.
<http://www.nber.org/papers/w22476>
- Philippon, T. (2012). Finance vs. Wal-Mart: Why are financial services so expensive? In A. Blinder, A. Lo, and R. Solow (Eds.), *Lessons from the Crisis*.
- Purda, L., and D. Skillicorn (2015) Accounting variables, deception, and a bag of words: Assessing the tools of fraud detection, *Contemporary Accounting Research* 32, 1193.
- Renault, T. (2017). Intraday online investor sentiment and return patterns in the US stock market. *Journal of Banking & Finance*, 84, 25-40.
- Rohner, Philippe and Uhl, Matthias, Robo-Advisors vs. Traditional Investment Advisors - An Unequal Game (2017). *Journal of Wealth Management*, Forthcoming.
- Rosenberg, Gili and Haghnegahdar, Poya and Goddard, Phil and Carr, Peter and Wu, Kesheng and Lopez de Prado, Marcos, (2015) Solving the Optimal Trading Trajectory Problem Using a Quantum Annealer. IEEE Journal of Selected Topics in Signal Processing, NYU Tandon Research Paper No. 2649376.
- Routledge, B. R., S. Sacchetto, and N. A. Smith (2018). Predicting merger targets and acquirers from text, Carnegie Mellon University working paper.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1985). *Learning internal representations by error propagation* (No. ICS-8506). California Univ San Diego La Jolla Inst for Cognitive Science.
- Russell, S., & Norvig, P. (1995). *Artificial intelligence a modern approach*. Artificial Intelligence. Pearson.
- Salchenberger, L. M., Cinar, E., & Lash, N. A. (1992). Neural networks: A new tool for predicting thrift failures. *Decision Sciences*, 23(4), 899-916.
- Samuels, A.L. (1959) Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3(3), 210-229.
- Shannon, C.E. (1940) A Symbolic Analysis of Relay and Switching Circuits, Master's Thesis, Department of Electrical Engineering, Massachusetts Institute of Technology.
- Sinha, N. R. (2010). Underreaction to news in the US stock market. Available at SSRN:
<http://ssrn.com/abstract=1572614>.

- Tam, K. Y., & Kiang, M. Y. (1992). Managerial applications of neural networks: the case of bank failure predictions. *Management science*, 38(7), 926-947.
- Tam, K. Y., & Kiang, M. (1990). Predicting bank failures: A neural network approach. *Applied Artificial Intelligence an International Journal*, 4(4), 265-282.
- Tang, J., & Yin, J. (2005). Developing an intelligent data discriminating system of anti-money laundering based on SVM. In *Machine Learning and Cybernetics, 2005. Proceedings of 2005 International Conference on* (Vol. 6, pp. 3453-3457). IEEE.
- Tetlock, P. C., Saar-Tsechansky, M., & Macskassy, S. (2008). More than words: Quantifying language to measure firms' fundamentals. *The Journal of Finance*, 63(3), 1437-1467.
- Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, 62(3), 1139-1168.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267-288.
- Turing, A (1950) Computing machinery and intelligence, *Mind*, 59(236), 433.
- Utans, J., & Moody, J. (1991). Selecting neural network architectures via the prediction risk: Application to corporate bond rating prediction. In *Artificial Intelligence Applications on Wall Street, 1991. Proceedings., First International Conference on* (pp. 35-41). IEEE.
- van Liebergen, B. (2017). Machine learning: A revolution in risk management and compliance? *Journal of Financial Transformation*, 45, 60-67.
- Vapnik, V. N. (1999). An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5), 988-999.
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3-28.
- Villalobos, M. A., & Silva, E. (2017). A Statistical and Machine Learning Model to Detect Money Laundering: An Application.
- Wall, L.D. (2017) Some Financial Regulatory Implications of Artificial Intelligence. Federal Reserve Bank of Atlanta Working Paper. Accessed at: <https://www.frbatlanta.org/-/media/documents/news/conferences/2017/1102-financial-regulation-fit-for-the-future/wall.pdf>.
- Williams, R. J. (1985) Learning Internal Representatives by Error Propagation, Institute for Cognitive Science Report 8506, San Diego, University of California.
- Winograd, T. (1971). Procedures as a representation for data in a computer program for understanding natural language (No. MAC-TR-84). MIT Project MAC.
- Wuermaling, J. (2018) Artificial Intelligence in Finance: Six warnings from a Central Banker. Intervention at the 2nd Annual Fintech Conference. Brussels.

Yadav, Y. (2015) How algorithmic trading undermines efficiency in capital markets, *Vanderbilt Law Review*, 68, 1607-1671.

Yeh, C. C., Chi, D. J., & Lin, Y. R. (2014). Going-concern prediction using hybrid random forests and rough set approach. *Information Sciences*, 254, 98-110.

Zhang, Y. and Trubey, P. (2018) Machine Learning and Sampling Scheme: An Empirical Study of Money Laundering Detection. Available at SSRN: <https://ssrn.com/abstract>

Endnotes

¹ Blockchain is an example of a distributed ledger technology that allows electronic payments to be made without going through a financial intermediary.

² A new Fintech sector, combining the real estate market and technology

³ “Robots in Finance Bring New Risks to Stability, Regulators Warn. Silla Brush,” Bloomberg Magazine. 1 Nov 2017. Available at: <https://www.bloomberg.com/news/articles/2017-11-01/robots-in-finance-bring-new-risks-to-stability-regulators-warn>

⁴ Asymptotic statistical properties stem from large sample theory and serve as a framework for statistical tests and estimated properties.

⁵ Or in other words...lots and lots of decisions trees.

⁶ https://en.oxforddictionaries.com/definition/artificial_intelligence

⁷ Renaissance Technologies would eventually become famous for its financial signal processing techniques for use in pattern recognition.

⁸ Fair Isaac Corporation

⁹ Sponsored by The US Department of Treasury.

¹⁰ “17 minute trading glitch put Goldman’s reputation on the line,” Arash Massoudi and Tracy Alloway, Financial Times. 22 August 2013. Available at: <https://www.ft.com/content/37fff9c6-0b36-11e3-bffc-00144feabdc0>

¹¹ “Britain Urged to Take Ethical Advantage in Artificial Intelligence,” John Thornhill, Financial Times. 16 April 2018. Available at: <https://www.ft.com/content/b21d1fb8-3f3e-11e8-b9f9-de94fa33a81e>

¹² “BlackRock bulks up research into artificial intelligence” Robin Wigglesworth and Chris Flood, Financial Times. 19 February 2018. Available at: <https://www.ft.com/content/4f5720ce-1552-11e8-9376-4a6390addb44>

¹³ With thanks to Mabel Lu of SWIFT for discussions

¹⁴ “China May Own More Artificial Intelligence Patents Than US By Year-End,” Pan Yue, China Money Network. 14 September 2017. Available at: <https://www.chinamoneynetwork.com/2017/09/14/china-may-hold-artificial-intelligence-patents-us-year-end>

-
- ¹⁵ “China Challenges the US for Artificial Intelligence Dominance,” Michael Woolridge, Financial Times. 15 March 2018. Available at: <https://www.ft.com/content/b799cb04-2787-11e8-9274-2b13fccdc744>
- ¹⁶ “China takes the crown in AI Funding,” Louise Lucas, Financial Times. 20 February 2018. Available at: <https://www.ft.com/content/1f58b248-1642-11e8-9376-4a6390addb44>
- ¹⁷ “The chip leap forward,” Will Knight, MIT Technology Review, Vol 122, 1. Jan/Feb 2019.
- ¹⁸ “Chinese AI Chipmaker Horizon Robotics raises up to \$1 billion,” Louise Lucas, Financial Times, 26 November 2018. Available at: <https://www.ft.com/content/bea3090c-f15c-11e8-ae55-df4bf40f9d0d>
- ¹⁹ “The chip leap forward,” Will Knight, MIT Technology Review, Vol 122, 1. Jan/Feb 2019.
- ²⁰ “China May Own More Artificial Intelligence Patents Than US By Year-End,” Pan Yue, China Money Network. 14 September 2017. Available at: <https://www.chinamoneynetwork.com/2017/09/14/china-may-hold-artificial-intelligence-patents-us-year-end>
- ²¹ “Beijing wants A.I. to Be Made in China by 2030,” Paul Mozur, New York Times. 20 July 2017. Available at: <https://www.nytimes.com/2017/07/20/business/china-artificial-intelligence.html>
- ²² A GLOBAL BRITAIN: From local startups to international markets. Available at: <http://coadec.com/Coadec-Report-A-Global-Britain.pdf>
- ²³ “BlackRock cofounder: Artificial intelligence won’t replace humans,” Oscar Williams-Grut, Business Insider. 3 November 2017. Available at: <http://www.businessinsider.com/blackrock-rob-kapito-tech-finance-artificial-intelligence-2017-11>
- ²⁴ “Wall Street’s King Quant David Shaw’s Secret Formulas Pile Up Money. Now He Wants a Piece of the Net,” James Aley, Fortune. 5 February 1996.
- ²⁵ “HSBC brings in AI to help spot money laundering,” Martin Arnold, Financial Times. 8 April 2018. Available at: <https://www.ft.com/content/b9d7daa6-3983-11e8-8b98-2f31af407cc8>
- ²⁶ “NatWest Uses Machine Learning to Fight Fraud,” Peter Walker. FsTech. 4 March 2018. Available at: http://www.fstech.co.uk/fst/NatWest_Machine_Learning_Fraud.php
- ²⁷ “NatWest Uses Machine Learning to Fight Fraud,” Peter Walker. FsTech. 4 March 2018. Available at: http://www.fstech.co.uk/fst/NatWest_Machine_Learning_Fraud.php
- ²⁸ “HSBC brings in AI to help spot money laundering,” Martin Arnold, Financial Times. 8 April 2018. Available at: <https://www.ft.com/content/b9d7daa6-3983-11e8-8b98-2f31af407cc8>
- ²⁹ Erica is a play on the bank’s name

³⁰ “Megabanks in Japan Embrace Artificial Intelligence,” Robot Technology. 30 October 2017. Available at: <https://business.inquirer.net/239571/megabanks-japan-embrace-artificial-intelligence-robot-technology>

³¹ “Allstate’s ‘Digital Colleague’ Amelia Answers Questions for Call Center Reps,” Sara Castellanos, *Wall Street Journal*. 30 March 2018. Available at: <https://blogs.wsj.com/cio/2018/03/30/allstates-digital-colleague-amelia-answers-questions-for-call-center-reps/>

³² “China’s debt collectors focus in on \$200 billion P2P debt pile,” Don Weinland, Financial Times. 5 June 2018. Available at: <https://www.ft.com/content/a219cd72-67a3-11e8-8cf3-0c230fa67aec>

³³ “Robo advisors start to Take Hold in Europe,” Gerrard Cowan, *Wall Street Journal*. 4 February 2018. Available at: <https://www.wsj.com/articles/robo-advisers-start-to-take-hold-in-europe-1517799781>

³⁴ “Banks look to the stars to spot trading mistakes,” Laura Noonan, Financial Times. 18 April 2018. Available at: <https://www.ft.com/content/578c2f26-430d-11e8-803a-295c97e6fd0b>

³⁵ DL should not be confused with high frequency trading (HFT) because DL is not looking to front-run trades or make money from speed of action.

³⁶ Kensho’s algorithm was originally called Warren after Warren Buffet

³⁷ “ING launches artificial intelligence bond trading tool Katana,” Laura Noonan, Financial Times. 12 December 2017. Available at: <https://www.ft.com/content/1c63c498-de79-11e7-a8a4-0a1e63a52f9c>

³⁸ Volatility refers to asset price fluctuations, or the standard deviation of price changes. To measure volatility with certainty can only be done ex-post. The term “trading volatility” means to speculate on future volatility, which is measured by implied volatility, which can be extracted from options prices for example.

³⁹ “Artificial intelligence is guiding venture capital to start-ups,” Maija Palmer, Financial Times. 11 December 2017. Available at: <https://www.ft.com/content/dd7fa798-bfcd-11e7-823b-ed31693349d3>

⁴⁰ To understand the challenges of black box algorithms, see articles such as <https://www.scientificamerican.com/article/demystifying-the-black-box-that-is-ai/>

⁴¹ The Flash Crash lasted half an hour during a trading. On 6 May the DJIA experienced its biggest one-day point decline (intraday basis) in its entire history. The stock prices of some of the world’s largest companies traded at incomprehensible prices: For example, Accenture traded at a penny a share, while Apple traded at \$100,000 per share (Kirilenko and Lo (2013)).

⁴² “UK pound plunges more than 6% in mysterious flash crash,” Jethro Mullen, CNN Business. 7 October 2016. Available at: <https://money.cnn.com/2016/10/06/investing/pound-flash-crash-currency-brexit/index.html>

⁴³ "Volatility how 'algos' changed the rhythm of the market," Robin Wigglesworth, Financial Times. 8 January 2019. Available at: <https://www.ft.com/content/fdc1c064-1142-11e9-a581-4ff78404524e>

⁴⁴ "Britain Urged to Take Ethical Advantage in Artificial Intelligence," John Thornhill, Financial Times. 16 April 2018. Available at: <https://www.ft.com/content/b21d1fb8-3f3e-11e8-b9f9-de94fa33a81e>

⁴⁵ <https://www.leverton.ai/>

⁴⁶ "How drones, data and AI are changing the property sector," Anthony Couse, World Economic Forum. 24 January 2018. Available at: <https://www.weforum.org/agenda/2018/01/proptech-drones-data-ai-property-sector/>

⁴⁷ "Hank Greenberg backs funding round for underwriting start up," Oliver Ralph, Financial Times. 6 December 2017. Available at: <https://www.ft.com/content/7d9867c8-d9c4-11e7-a039-c64b1c09b482>

⁴⁸ "Hong Kong stock exchange enlists AI in fight against rule breakers," Emma Dunkley, Financial Times. 16 April 2018. Available at: <https://www.ft.com/content/64bcb136-412b-11e8-803a-295c97e6fd0b>

⁴⁹ "Bank of England Experiments with Artificial Intelligence and Blockchain Tech," Brady Dale, The Observer. 12 July 2017. Available at: <https://observer.com/2017/07/bank-of-england-mindbridge-ripple/>


⁵⁰ Mifid II or Markets in Financial Instruments Directive came into effect early 2018 and is designed to offer greater protection for investors and inject more transparency into all asset classes: from equities to fixed income, exchange traded funds and foreign exchange.

⁵¹ Serious Fraud Office CEO Ben Denison reveals how AI is transforming legal work, Thomas Macaulay, CIO, 3 January 2018.

⁵² "Artificial intelligence could yet upend the laws of finance," Stuart Kirk, Financial Times. 22 January 2018. Available at: <https://www.ft.com/content/8c263c06-fc70-11e7-9b32-d7d59aace167>

⁵³ "Artificial intelligence could yet upend the laws of finance," Stuart Kirk, Financial Times. 22 January 2018. Available at: <https://www.ft.com/content/8c263c06-fc70-11e7-9b32-d7d59aace167>

⁵⁴ As an example, Finland has rolled out free-access university courses to enable awareness on basic concepts of AI technology. In 2017, Finland was the first EU country to put a national AI strategy into writing.



turing.ac.uk
@turinginst