The impact of machine learning on UK financial services

Bonnie G. Buchanan* and Danika Wright**

Abstract: Machine learning is an increasingly key influence on the financial services industry. In this paper, we review the roles and impact of machine learning (ML) and artificial intelligence (AI) on the UK financial services industry. We survey the current AI/ML landscape in the UK. ML has had a considerable impact in the areas of fraud and compliance, credit scoring, financial distress prediction, robo-advising and algorithmic trading. We examine these applications using UK examples. We also review the importance of regulation and governance in ML applications to financial services. Finally, we assess the performance of ML during the Covid-19 pandemic and conclude with directions for future research.

Keywords: machine learning, AI, financial services, big data

JEL classification: C45, G0, G1, G2, G15, G 20, O30

Finance is really a wonderfully pure information-processing business.

D. E. Shaw, 1996

Artificial Intelligence (AI) and machine learning (ML) techniques . . . [have] the potential to yield enormous benefits for households and businesses by opening up new lines of credit, providing greater choice, better-targeted products and keener pricing.¹

Mark Carney,
Governor of Bank of England, 2019

I. Introduction

The Fourth Industrial Revolution is thought of as a series of transformative technological breakthroughs in areas including artificial intelligence (hereafter AI) and machine learning (hereafter ML) that are currently disrupting and innovating businesses as well

* Surrey Business School, University of Surrey, UK, e-mail: b.buchanan@surrey.ac.uk

** Discipline of Finance, University of Sydney, Australia, e-mail: danika.wright@sydney.edu.au

The authors wish to thank participants at the OxREP workshop and the anonymous reviewers. They also wishes to thank David Bholat, Daniel Susskind, and Alison Gomm. Buchanan wishes to thank participants at the March 2020 Bank of England Workshop on AI in Financial Services.

¹ 'Enable, Empower, Ensure: A New Finance for the New Economy', Mark Carney, Mansion House speech, 20 June 2019.

<https://doi.org/10.1093/oxrep/grab016>

© The Author(s) 2021. Published by Oxford University Press.

For permissions please e-mail: journals.permissions@oup.com

as challenging traditional policy-making approaches (Schwab, 2017). Bholat (2020) attributes the swift rise in AI to the variety, volume, and velocity of data created in the Internet era. While uptake of AI in the financial services industry took some time, AI now (and especially ML) dominates many lists of key trends in financial services and Fintech (Buchanan, 2019; CB Insights, 2019). In a 2019 PwC survey, 56 per cent of global executives surveyed considered AI to be a strategy to transform financial services. Jamie Dimon, CEO of J.P. Morgan, states that using AI in fraud detection alone saves his bank \$150 million per year.² In 2018, the UK government announced plans to spend \$1.3 billion on an AI push.³ AI could potentially add an additional US\$814 billion (£630 billion) to the UK economy by 2035, with the financial services industry an area for substantial activity (Hall and Pesenti, 2017).

During the last few years, adoption of AI and ML has been particularly widespread in the financial services sector. In the UK a new AI startup has been established on roughly a weekly basis since 2014.⁴ London is a main hub for AI startups, and some of the prominent firms include Swiftkey, DeepMind, and Ravn. According to a 2018 CognitionX report, London's role as an AI supplier⁵ base is twice the size of Berlin and Paris combined. London is also extremely well poised to be an AI and ML leader in finance and insurance, with new supplier formation growing at an annual rate of 42 per cent (compared with the global rate of 24 per cent annually; CognitionX (2018)). The report attributes the rapid AI growth in London to several factors: a large, multicultural, technical and entrepreneurial talent pool, access to leading financial institutions and markets, a strong peer-support network, and increasing international investment.

The history of AI development is well-documented (Buchanan, 2019; Wooldridge, 2020; Russell, 2021) and is summarized in Appendix I. John McCarthy coined the term 'artificial intelligence' in 1953. There are a variety of AI definitions, including Kaplan's (2016) succinct description that,

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 conclusions are drawn with minimal information, the more intelligent the behaviour. (2016, pp. 5–6)

In this paper, we follow the Financial Stability Board (FSB) definition of AI as, 'the theory and development of computer systems able to perform tasks that have traditionally required human intelligence' (FSB, 2017, p. 4).

The origins of ML are attributed to McCulloch and Pitts (1943). Considered to be a subset of AI, ML uses algorithms to automatically optimize through experience with limited or no human intervention. ML is primarily derived from sources such as experience, practice, training, and reasoning (Buchanan, 2019). ML techniques include general pattern recognition and universal approximations of observable patterns in data where no *a priori* analytical solution exists (Cybenko, 1989). McCulloch and Pitts

² <https://www.ft.com/content/11aab1cc-907b-11ea-bc44-dbf6756c871a>

³ <https://www.gov.uk/government/news/13-billion-industry-government-investment-in-uk-economy-and-new-partnership-driving-early-disease-detection>

⁴ <http://coadec.com/Coadec-Report-A-Global-Britain.pdf>

⁵ An AI supplier is one who sells at least one AI product whether or not they sell non-AI products.

(1943) recognized that brain signals are digital in nature, specifically displaying binary signals. According to [Chakraborty and Joseph \(2017\)](#), ML systems comprise five components: (i) a problem, (ii) a data source, (iii) a model, (iv) an optimization algorithm, and (v) validation and testing. The ideal situations for ML are those that require extracting patterns from noisy data or sensory perception. The four main drivers of the growth in ML include: (i) the transition from physical to electronically stored data, (ii) improvements in memory and computing speed, (iii) easier access to data due to the Internet, and (iv) low-cost high-resolution digital sensors ([Buchanan, 2019](#)).

A subset of ML is deep learning (DL), which [FSB \(2017\)](#) defines as

a form of machine learning that uses algorithms that work in ‘layers’ inspired by the structure and function of the brain. Deep learning algorithms, whose structure are called artificial neural networks, can be used for supervised, unsupervised, or reinforcement learning.

DL is considered to be a statistical technique for finding patterns in large amounts of data. Much of DL is underpinned by neural networks which mimic the way multiple layers of the brains’ *neurons* work (hence the term ‘deep’). In summary, DL is a subset of ML, which in turn is a subset of AI.

Financial services ML can be categorized according to the following broad classifications: (i) fraud detection and compliance; (ii) credit risk analysis and prediction of financial distress; (iii) banking chatbots and robo-advisory services; and (iv) algorithmic trading.

Each of these categories has had a considerable impact on the UK financial services sector, and the outlook is optimistic ([Hall and Pesenti, 2017](#); [CognitionX, 2018](#)). A 2019 Bank of England (BoE) survey finds that two-thirds of financial institutions surveyed currently use some form of ML, with the majority of uses in banking and insurance ([Bank of England, 2019](#)).

In this paper, we examine trends in the application of AI and ML in the UK financial services industry and the overall impact. Moreover, we discuss challenges and opportunities facing the sector alongside advances in AI and ML use. We also detail how ML and AI are applied to big data and how algorithms are used to generate value in the commodification of data in financial services.

The paper is set out as follows: in the next section we discuss the current AI/ML landscape, with an emphasis on the UK. Machine learning applications in UK financial services follow, with more detailed discussions on fraud detection, compliance, credit scoring, financial distress prediction, robo-advisory services, and algorithmic trading. We then discuss the importance of regulation and ML. We then discuss the impact of quantum computing on ML development in the UK. The importance of good governance in AI/ML follows and the final section concludes.

II. The current AI/ML landscape in the UK

The UK ranks third globally in terms of the number of AI companies based there, coming after China and the USA ([Buchanan, 2019](#)). China’s rapid rise in ML is attributed to the massive supply of data (for example, WeChat has over a billion users) and

Figure 1: Importance of big data analytics and ML technologies worldwide, 2018

Source: Statista.

the growth of facial recognition and AI chips in the country. The increase in AI chip revenue has the potential of boosting economic growth and productivity, improving societal growth, and generating savings. The UK is looking to also capitalize on this—especially in the field of financial services. In Figure 1 we observe that Spark MLlib (a machine learning library) and TensorFlow (a machine learning platform) are rated the most critically important technologies worldwide.

In terms of AI's contribution to the UK economy, Hall and Pesenti (2017) estimate AI could add an additional US\$814 billion (£630 billion) by 2035. Since 2014⁶ new AI startups have rapidly established in the UK. Swiftkey, DeepMind, Ravelin, Featurespace, Previs, and Crowdemotion are among the more prominent UK AI companies.

In 2015 Blackrock acquired automated investment platform FutureAdvisor and invested £60 million into early-stage UK venture capital fund Forward Partners, which targets applied AI startups.⁷ In January 2020, the UK had the fourth highest concentration of unicorns globally.⁸ Ten UK headquartered unicorns were active in the finance, insurance, and real estate industry, ahead of the four unicorns specifically dedicated to health and pharmaceuticals (Statista, 2020).

Table 1 displays the number of AI firms and AI-Fintech firms in 2020 by country according to region. Firms are categorized as belonging to either AI or Fintech, based on the Crunchbase category groups. Table 1 shows that the highest number of AI firms and AI-Fintech firms are located in the USA, followed by the UK, India, and Israel.

⁶ <http://coadec.com/Coadec-Report-A-Global-Britain.pdf>

⁷ <http://www.businessinsider.com/blackrock-rob-kapito-tech-finance-artificial-intelligence-2017-11?r=UK&IR=T>

⁸ A 'unicorn' is a company with a valuation of at least a \$1 billion (Buchanan and Cao, 2018).

China, Canada, and Germany also have a high number of AI firms. India's national AI strategy has focused on the social sector (i.e. health, education, agriculture, smart cities, and smart mobility) and AI investment leapt from US\$44 million in 2016 to US\$77 million in 2017.⁹ The Israeli AI sector raised close to \$2 billion in 2018 and is growing rapidly. Many Israeli AI entrepreneurs have a strong military background, and as a result have more extensive hands-on experience working with AI, image processing, and data science than many other international entrepreneurs. The main AI technologies in Israel include machine learning, deep learning, computer vision, natural language processing, robotics, and speech recognition.¹⁰

Table 2 provides a more granular comparison of the UK and US AI and Fintech markets. The headquarters for AI and Fintech firms is tabulated for both countries. We observe that UK AI and Fintech firms are much more geographically concentrated than those in the US. Most UK AI firms are headquartered in the south of the nation, with London accounting for nearly three-quarters of AI headquarters, followed by Cambridge. Fintech firms are concentrated in London (74.7 per cent) followed by Manchester (1.6 per cent). In the US, the picture is very different. California is the preferred headquarter location for both AI and Fintech firms (44.5 per cent and 29.2 per cent, respectively). New York is the second most popular location (with 14.2 per cent and 17.2 per cent, respectively). California is the preferred choice due to proximity to Silicon Valley firms, and New York also has the advantage of proximity to Wall Street firms.

In Table 3 we present the 2020 world rankings, based on valuation, of the largest leading UK AI start-ups. Graphcore, based in Bristol, is ranked 15th globally and is a unicorn chipmaker valued at \$1.5 billion.¹¹ Incorporated in 2016, Graphcore develops intelligence processing units which accelerate ML applications. Its AI chips run ML applications in industries including health, automotive, manufacturing, security, and financial services. For example, hedge fund Carmot Capital experienced a 26 times increase in speed of its applications due to Graphcore chips. Its intelligence processing chips have a competitive advantage in voice processing, image recognition, and video analysis.¹² Darktrace is an AI firm specializing in cybersecurity, and is backed by investment vehicle Invoke Capital, KKR, and Softbank. Its clients include the NHS, Drax (the UK's biggest power station), and Gatwick airport.¹³ Benevolent AI specializes in ML tools for driving drug discovery and development. At present, it is using ML techniques to search for treatments for the coronavirus. Onfido, founded in 2012, uses ML technology in identity recognition and fraud protection. Babylon Health, founded in 2013, incorporates ML technology to its chatbot function which is used by the NHS. Finally, UK Tractable, ranked 96th, is active in the insurance industry. Damage and

⁹ <https://tracxn.com/reports/fMfLAJmqwTD4z61mrkRwbUsWTnIETKriWGCELb>

¹⁰ 'The Thriving AI Landscape in Israel and What It Means for Global AI Competition', Gil Press, *Forbes*, 24 September 2018. Available at: <https://www.forbes.com/sites/gilpress/2018/09/24/the-thriving-ai-landscape-in-israel-and-what-it-means-for-global-ai-competition/?sh=10f683ae30c5>

¹¹ 'Microsoft signs cloud AI deal with UK chipmaker Graphcore', Madhumita Murgia, *Financial Times*, 13 November 2019.

¹² Ibid.

¹³ 'Inside Darktrace, the UK's \$1.65 billion cybersecurity start-up', Aliya Ram, *Financial Times*, 10 October 2018.

Table 1: Number of AI firms and Fintech–AI firms by Region, 2020

Asia-Pacific			Americas			Europe			Middle East and Africa		
Country	AI firms	AI-FT firms	Country	AI firms	AI-FT firms	Country	AI firms	AI-FT firms	Country	AI firms	AI-FT firms
IND	545	43	USA	5,093	486	GBR	1,052	156	ISR	358	28
CHN	429	24	CAN	529	51	DEU	434	33	ZAF	28	2
AUS	166	19	BRA	240	19	CHE	158	16	ARE	55	5
SGP	198	34	MEX	55	5	ESP	199	15	NGA	16	3
HKG	97	18	ARG	29	2	FRA	347	37	KEN	18	6
JPN	116	7	CHL	33	4	NLD	173	12	EGY	15	2
IDN	24	1	COM	19	6	IRL	75	6	GHA	4	1
MYS	35	8	CYM	4	1	SWE	115	14	UGA	3	0
KOR	81	5	PER	6	0	RUS	67	3	LBN	11	2
PHL	11	0	ECU	2	1	ITA	119	14	SAU	4	0
NZL	22	3	URY	5	0	EST	36	9	IRN	11	0
THA	14	1	TTO	2	0	DNK	67	6	KWT	1	0
VNM	15	2	PRI	2	0	BEL	64	9	TUN	9	1
TWN	56	5	CRI	2	0	POL	102	11	CMR	3	0
PAK	8	0	VEN	2	1	FIN	84	7	MAR	4	0
BGD	9	0	DOM	1	0	TUR	49	5	JOR	2	0
LKA	3	0	Region	6,024	576	AUT	67	4	ARM	3	1
KHM	1	0				NOR	46	9	AZE	1	0
MMR	1	0				LTU	15	0	ZMB	1	0
NPL	3	0				PRT	54	5	BHR	3	0
KAZ	2	0				MLT	9	3	Region	550	51
Region	1,836	170				CYP	13	4			
						LUX	11	3			
						ROM	26	2			
						UKR	36	2			
						BGR	17	0			
						CZE	28	6			
						GRC	23	3			
						HUN	27	2			
						LVA	16	1			
						GIB	2	1			
						SVN	12	2			
						BLR	12	1			
						ISL	12	0			
						HRV	2	0			
						SVK	8	0			
						IMN	1	0			
						GEO	2	0			
						SRB	8	1			
						Region	3,588	402			

Source: Crunchbase.

repair costs are estimated in real time, using ML visual analysis of images of the damaged item.

External investment funding is a key input to AI firm and industry growth. [Figure 2](#) charts the average funding per firm by year into AI and Fintech for the UK and US since 2000. We observe that UK Fintech firms, on average, received more investment funding throughout the 2007–8 financial crisis than the average US Fintech firm. This is likely due to the surge in investment with peer-to-peer funding which started in the UK

Table 2: Top 20 regions for AI and Fintech firms' HQs for UK and US

United Kingdom				USA			
AI firms		Fintech firms		AI firms		Fintech firms	
London	72.5%	London	74.7%	California	44.5%	California	29.2%
Cambridge	3.3%	Manchester	1.6%	New York	14.2%	New York	17.2%
Oxfordshire	2.8%	Oxfordshire	1.4%	Massachusetts	6.0%	Texas	5.7%
Edinburgh,	1.6%	Kent	1.2%	Texas	5.0%	Florida	5.1%
Manchester	1.5%	Bristol, C	1.1%	Washington	2.8%	Illinois	4.4%
Bristol, C	1.2%	Surrey	1.0%	Illinois	2.5%	Massachusetts	4.3%
Surrey	1.1%	Edinburgh,	0.9%	Florida	2.5%	Georgia	2.7%
Leeds	0.8%	Cambridge	0.8%	Virginia	1.9%	Washington	2.6%
Newcastle	0.8%	Glasgow	0.7%	Georgia	1.8%	Pennsylvania	2.2%
Glasgow	0.7%	Hertford	0.7%	Pennsylvania	1.8%	New Jersey	2.1%
Birmingham	0.7%	Essex	0.6%	Colorado	1.8%	Colorado	2.1%
Belfast	0.7%	Cheshire	0.6%	Maryland	1.2%	Virginia	1.7%
Essex	0.6%	Birmingham	0.5%	New Jersey	1.2%	North Caro	1.6%
Hampshire	0.6%	Hampshire	0.5%	North Caro	1.2%	Arizona	1.4%
East Sussex	0.6%	Nottingham	0.5%	Ohio	0.9%	Utah	1.3%
Nottingham	0.5%	Leeds	0.4%	District o	0.9%	Ohio	1.3%
Buckingham	0.5%	Sheffield	0.4%	Michigan	0.7%	Michigan	1.2%
Cardiff	0.5%	East Sussex	0.4%	Oregon	0.7%	Connecticut	1.2%
Sheffield	0.4%	Buckingham	0.4%	Utah	0.7%	Nevada	1.1%
Norfolk	0.4%	Bath and N	0.4%	Delaware	0.7%	Maryland	1.1%

Source: Crunchbase.

Table 3: World rankings of UK AI start-ups, 2020

UK AI start-up	World ranking
Graphcore	#15
Darktrace	#19
Benevolent AI	#24
Babylon Health	#54
Onfido	#68
Tractable	#96

Source: Statista.

at this time. US AI funding dominates UK AI funding for most of the sample period, although the UK does exceed US AI funding in 2007, 2008, 2011, 2013, and 2015.

In [Figure 3](#) we compare the number of ML firms founded in China, the US, and the UK. The US clearly leads in number of ML firms founded, followed by the UK and China. For both the US and UK, growth accelerated post-financial crisis.

[Figure 4](#) examines growth in DL startups since 2008 (a shorter period due to the more recent emergence of this category of AI specialization). Panel A details how DL growth escalates after 2010 and is dominated by the US, followed by the UK and China. In Panel B we observe that the share of DL firms based in the UK is increasing compared to the rest of the world. The proportion of UK DL start-ups more than doubles between 2018 and 2019.

In terms of employee numbers, firms in AI industries are usually fairly small. Using employee count to measure firm size seems a more reasonable approach than revenue or earnings, as many AI start-up firms have low/volatile financial returns in early years.

Figure 2: Funding into AI and Fintech firms in UK and US, 2000–19**A: AI Funding** **B: Fintech funding**

Note: These figures are standardized in US dollars.

Source: Crunchbase.

We see in [Figure 5](#) that despite the larger market in the US, the distribution of firm sizes is comparable to UK AI market.

III. Machine learning applications in UK financial services

Personalized financial planning, fraud detection, anti-money-laundering, and process automation are key areas identified by [Hall and Pesenti \(2017\)](#) that have great potential in the AI sector. The financial services industry is the biggest spender on AI services

Figure 3: Top 3 ML by firms founded

Source: Crunchbase.

outside of the technology sector (Citi, 2018). In 2019, the Bank of England (BoE) and Financial Conduct Authority (FCA) surveyed 300 financial market participants (including banks, trading platforms, and fund managers). Of the respondents,¹⁴ 66 per cent said they currently employ ML and often in several areas (Bank of England, 2019). According to respondents, ML is used in both front- and back-office settings, fraud prevention, anti-money-laundering (AML), and customer service situations. The median firm surveyed by BoE/FCA uses ML applications in two company areas, and this adoption is expected to double over the next 3 years. The survey finds that many firms are now deploying ML in advanced stages, with the most prominent in the banking and insurance areas.

Until recently, hedge funds and high-frequency trading firms were the main users of AI in finance, but applications have now spread to other areas, including banks, regulators, Fintech, and insurance firms. Within the financial services industry, AI applications include algorithmic trading, credit risk assessment, wealth management, debt collection, cyber risk, portfolio construction and optimization, model validation and back-testing, robo-advising, virtual customer assistants, market impact analysis,

¹⁴ The survey received 106 responses out of 300 surveyed.

Figure 4: Panel A: Top 3 DL, by firms founded **Panel B:** DL firms started in UK, % of global

Source: Crunchbase.

regulatory compliance, and stress testing. In this section, we discuss four specific areas in which AI/ML is currently changing the financial services industry, namely (i) fraud detection and compliance; (ii) credit risk analysis and prediction of financial distress; (iii) banking chatbots and robo-advisory services; and (iv) algorithmic trading.

Figure 5: Firm size: UK vs US

Source: Crunchbase.

(i) Fraud detection and compliance

Cybercrime costs the global economy over \$600 billion annually, and credit card fraud accounts for most of this (Statista, 2020). In the same report, 69 per cent of CEOs in financial services state they are extremely concerned about cybercrime (compared with 61 per cent across other sectors). According to the FCA, UK banks spend £5 billion a year combating financial crime.¹⁵ In the UK between 2015 and 2016, there was a 66 per cent increase in the number of reported cases of payments-related fraud.¹⁶ Substantial fines have been imposed upon financial institutions for failing to stop financial fraud. As a result, banks have turned to AI/ML techniques to improve their operations and

¹⁵ 'HSBC brings in AI to help spot money laundering', Martin Arnold, *Financial Times*, 8 April 2018.

¹⁶ 'NatWest uses machine learning to fight fraud', Peter Walker, FsTech website, 4 March 2018.

regulatory compliance. For example, PayPal has used ML to reduce its fraud rate to just 0.32 per cent compared with the industry average of 1.32 per cent.¹⁷

However, fraud detection involves more than just a checklist of risk factors. Fraud is a latent variable and, compared to making accurate predictions of shopping decisions, it is more difficult to identify possible financial fraud. This is because retailers have access to full customer transaction histories (Bauguess, 2017; Buchanan, 2019). Prior to the widespread adoption of ML, a rules-based approach has been the preferred approach to AML monitoring. ML produces a more efficient, accurate approach to AML monitoring (Proudman, 2019). Fraud detection systems include ML techniques that actively learn and calibrate both real and potential security threat responses. Anomalous behaviours are then flagged for further investigation.

One of the most successful applications of ML is credit card fraud detection. Algorithm training, back-testing, and validation are based on very large datasets of historical credit card transaction data. 'Fraud' versus 'non-fraud' events are labelled by classification and fraudulent transactions can then be stopped in real time (van Liebergen, 2017).

Many UK financial institutions use ML fraud prevention alternatives. For example, NatWest adopted a ML solution, named Corporate Fraud Insights, 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.¹⁸ In fighting money laundering, HSBC uses software by Quantexa AI, a UK based start-up.¹⁹

Another challenge arises from declining transactions too aggressively in order to prevent fraud. Transactions wrongly declined due to suspected fraud are known as a 'false positive' and represent just as big a threat to the financial services industry as actual fraud. This disadvantages the issuer because a false-positive declined transaction can result in erosion of customer loyalty to the financial institution. ML methods can substantially reduce false declines and improve credit card approvals.

Supervised ML algorithms in fraud detection include Bayes logistic regression, decision tree, random forest, support vector machines, and artificial neural networks. Sentiment analysis, where ML learning attempts to discover new trends and signals in financial activity and then replicate and enhance human 'intuition' is also used in fraud detection. Support vector machines are also used for detecting corporate management fraud using basic financial data (Cecchini *et al.*, 2010).

In money laundering detection cases, a large number of false positive results are created which in turn creates a number of challenges. First, there are few large public datasets that can be used to predict money laundering. Many financial institutions still rely on a traditional rules-based system that emphasizes individual transactions and simple transaction patterns. This may not be sophisticated enough to detect complex transactions (van Liebergen, 2017).

¹⁷ 'How PayPal boosts security with artificial intelligence', Michael Morisy, *MIT Technology Review*, 25 January 2016.

¹⁸ 'NatWest uses machine learning to fight fraud', Peter Walker, FsTech website, 4 March 2018.

¹⁹ 'HSBC brings in AI to help spot money laundering', Martin Arnold, *Financial Times*, 8 April 2018.

In response to the £1.2 billion money laundering fine imposed on HSBC in 2012, the bank adopted the Google cloud ML for use in its AML activities. In 2016, Lloyds Banking Group partnered with US-based AI startup Pindrop to combat unusual financial activity and fraud. Pindrop's ML technology is called 'phoneprinting', a system which analyses 147 unique human voice features to detect fraudulent activity, like a vocal fingerprint. Lloyds Banking Group was the first European organization to use 'phoneprinting' and it uses the technology across the Lloyds, Bank of Scotland, and Halifax networks. In the US, it has been used to detect over 80 per cent of fraudulent activity.²⁰ Incorporated in 1997, UK-based Eckoh is valued at £163 million²¹ and is quoted on the Alternative Investment Market. Eckoh uses interactive voice response as part of its audio tokenization for secure phone payments. It also uses text recognition for chat payments. For Eckoh, ML utilization has improved timeliness, data security, and efficiency as well as regulatory requirements, minimizing the potential for fraud. Privately held London-based SiloBreaker applies ML and textual analysis to identify security threats. SiloBreaker uses unstructured data extracts from across the internet, including the dark web. The data extracted covers blogs, social media, and news. In terms of output, SiloBreaker's output directly helps financial services firms. Benefits are measured by the real time insights of issues and developments detected in other industries or companies that give insight to trade opportunities.

The UK Serious Fraud Office (SFO) includes ML techniques in fraud and prosecution cases. On average, the SFO processes over 100 million documents annually. One of the SFO's most famous cases is the Rolls Royce bribery case, which resulted in the largest ever fine imposed in the UK for criminal conduct.²² The SFO used the RAVN robotic system, costing £50,000, saving 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. RAVN processed 30 million documents at a rate of up to 600,000 per day in the Rolls Royce case (compared with a team of lawyers that would have processed 3,000 per day) (Buchanan, 2019). The SFO also uses OpenText Accelerate to connect people during their fraud inquiries. Data from the Enron corporate scandal is used as the training set for the SFO ML algorithm.

(ii) Credit risk

Spikes in credit risk have been at the centre of previous financial crises. Credit risk is therefore a primary risk facing financial institutions. Credit risk analysis is used for mortgage, auto, credit card, and personal loans. Durand (1941) first introduced discriminant statistical analysis to distinguish between good and bad loans. In the financial services industry, most consumer loans are decided based on automated credit. Early uses of ML approaches in credit scoring include Makowski (1985) and Desai *et al.* (1996). The 2007–8 financial crisis revealed weaknesses in established credit risk

²⁰ <https://www.finextra.com/newsarticle/29574/lloyds-to-use-pindrop-tech-to-identify-fraudulent-calls>

²¹ Source: Capital IQ.

²² 'Serious Fraud Office CEO Ben Denison reveals how AI is transforming legal work', Thomas Macaulay, *CIO*, 3 January 2018.

management techniques. Credit scoring lends itself to different ML approaches, such as supervised learning to construct consumer credit risk profiles. Specific ML techniques to assess credit risk profiles include tree-based classifiers, random forest models, radial basis functions, and support vector machines.

Davis *et al.* (1992) apply two ML methods, specifically the Gamberman and Thatcher (1990) ‘G&T’ Bayesian inference model and neural networks, to Bank of Scotland credit card data. After the 1980s, logistic regression and linear programming were the mainstays of credit risk analysis (Thomas, 2000).²³ Neural networks and support vector machines (SVMs) complement traditional methods because they provide the capability to model curvature and non-linearity of data. SVMs are supervised learning algorithms that are applied in image recognition and text categorization and increase accuracy and credit scoring reliability (Bellotti and Crook, 2009).

A key benefit of ML methods in credit analysis is that their increased accuracy in predicting delinquencies and default forecasts reduces the likelihood of false positives, which in turn improves the value and authenticity of credit card debt securitization. Forecasts are more adaptive and can respond to the dynamics of changing credit cycles as well as the absolute levels of default rates (Khandani *et al.*, 2010). Khandani *et al.* (2010) use classification and regression trees (CART) to construct credit-risk forecasts. CART is a type of predictive algorithm in ML, and in this context puts input variables through a recursive sequence of simple binary relations to predict delinquencies and defaults. CART is able to detect the nonlinear interactions between input variables. Khandani *et al.* (2010) find that ML forecasts yield savings of between 6 and 25 per cent of total losses.

ML applications in credit market default prediction are gaining a wider audience. ML can be used to appraise character and reputation variables when assessing future payment behaviour. Using 250,000 purchases from a German e-commerce company, Berg *et al.* (2020) analyse ten digital footprints variables for default prediction. They find that digital footprints complement credit bureau information, affecting access to credit and reducing default rates. They find that digital footprints work equally well for un-scorable as well as for scorable customers. For customers with good digital footprints, default rates drop significantly after adoption of the digital footprint.

Logistic regression modelling is widely used in the mortgage literature. However, logistic regression misses the highly non-linear influence that variables such as unemployment have on borrower behaviour. A critical issue arises because the highly non-linear relationship between borrower behaviour and risk factors throws previous conclusions drawn from linear mortgage models into doubt. Sirignano *et al.* (2021) apply DL (specifically neural networks) to model mortgage credit and prepayment risk. Using a dataset of 120 million mortgages between 1995 and 2014, which represents 70 per cent of all US mortgages, they find a stronger relation between housing finance markets and the macroeconomy than previously thought. The output of their model represents conditional probabilities of different mortgage states (current, 30-days late, prepaid, foreclosed, etc.). The improved forecasts can help improve mortgage-backed security valuation, especially when facing adverse macroeconomic conditions. Findings can also have huge social implications as well. Along with logistic

²³ Thomas (2000) provides an excellent survey of the early days of credit scoring.

regressions, [Fuster *et al.* \(2020\)](#) use ML technology to examine the distributional implications of better statistical technology. They also use US mortgage data (10 million between 2009 and 2013) and find that ML methods produce improvements in predictive powers. [Fuster *et al.* \(2020\)](#) examine how risky and less risky borrowers are distributed across societally important categories such as race, income, and gender, finding increased disparities in certain consumer outcomes. Under ML predictions, white, non-Hispanic groups experience lower estimated default propensities and black and Hispanic borrowers face larger charges and are less likely to gain from the new technology.

London-based CreditEnable focuses on credit analytics and has obtained \$7.9 million in seed funding over five rounds.²⁴ This startup uses public financial and leading indicator data, coupled with proprietary data from lenders to assess creditworthiness and also build a proprietary database of SME borrowers. CreditEnable uses DL to create a managed marketplace that matches credit and risk profiles between lenders and SME borrowers. The startup has measured success through improved borrowing cost efficiency and timeliness.

(iii) Predicting financial distress

Credit scoring and bankruptcy prediction are the two main aspects in financial distress prediction. ML techniques (primarily neural networks) to predict bank distress were introduced in the 1990s. Most neural network approaches use a multilayer method. [Atiya \(2001\)](#) develops a neural network model for bankruptcy prediction that provides a marked improvement for out-of-sample forecasts. [Tang *et al.* \(2020\)](#) incorporate textual and management factors into ML-based financial distress predictions. SVMs are frequently chosen as the base learning model in credit scoring ensemble algorithms. [Suss and Treitel \(2019\)](#) use ML techniques to create an early warning for UK bank distress. They apply *k* nearest neighbours (kNN), random forests, boosting, and SVMs to a sample of UK banks between 2006 and 2012, finding random forest algorithms to be superior to the other ML techniques. Finally, sentiment analysis is shown to have beneficial use in financial distress modelling ([Wang *et al.*, 2018](#)).

In predicting financial distress, it is important to distinguish between a false positive (a Type I error) and a Type II signal. In a bankruptcy context, a Type I error predicts a firm will be fine, but it turns out the bank is in distress. On the other hand, a Type II error predicts a firm is going to be in distress, but it turns out to be fine. Clearly, the Type I error is the costlier mistake to the financial institution and exacerbates enterprise risk. It is also very serious and costly from a regulator's point of view. [Liang *et al.* \(2018\)](#) specifically address lowering the Type I error issue. They introduce a classifier ensemble principle, which is designed to reduce the occurrence of Type I errors. The classifier ensemble principle uses a 'divide and conquer' approach, whereby a complex problem is broken down into small sub-problems. A specific classification technique (such as SVM, NN, or decision tree) is then applied to each sub-problem. [Suss and Treitel \(2019\)](#) find ML methods are superior to classical approaches for predicting bank distress in the UK and that ML techniques can benefit from financial regulation.

²⁴ https://www.crunchbase.com/search/funding_rounds/field/organizations/funding_total/creditenable

(iv) Robo-advisory services

ML algorithms have enabled tools for individuals to access semi-personalized portfolio management services through automated online platforms, known as robo-advisors. A robo-advisor requests an individual to submit investment preference information which it uses to automatically create and manage the individual's investment portfolio. While the interface may include an automated chatbot, robo-advisors essentially remove the direct interaction between investor and human financial advisor. Portfolio optimization tools are used by human advisor and robo-advisor services alike and offer a much greater level of investment sophistication to the typical retail investor.

The key benefits for retail investors between robo-advisors and human advisors are in cost and accessibility. Robo-advisors are less expensive and can be accessed anytime from virtually any location. Greater accessibility also comes through the ability to use the service with minimum investment thresholds. Robo-advisors and chatbots are powered by natural language processing (NLP) and ML algorithms which have become powerful tools to provide a personalized, conversational, and natural experience across different domains. They have achieved significant appeal with millennial consumers who do not need a physical advisor to feel comfortable investing (Buchanan, 2019).

The advantages of robo-advisory services are: (i) passive market access with strategic asset allocation strategies, that often include more client customization than the traditional risk profiles used by banks, (ii) cost-efficient implementation of a diversified asset allocation, and (iii) a reduction of behavioural biases. Compared to traditional investment advice, robo-advisors can produce cost savings (assuming a traditional human advisor fee of 1.5 per cent) and improve net portfolio performance.

The growth of robo-advisory platforms increased rapidly in the Fintech startup space following the 2007–8 financial crisis. Currently the UK is a dominant player in the European robo-advisory industry. The recent low-interest-rate environment in the UK has been a major driver of this growth. As interest earned from bank deposits has become negligible, demand has grown for investment services where robo-advisors provide an inexpensive but convenient platform. Robo-advisors have the potential to lower costs and increase the quality and transparency of financial advice for consumers.

Robo-advisors have expanded to offer diverse products including pensions, mortgages, exchange traded funds (in stocks and bonds for portfolio construction), health insurance, and taxation services. UK financial firms are also examining ways to compete in a 'hybrid-robo' market, which combines the online digital platform with some limited human advisor interaction. For example, UK-based MoneyFarm is in this category and aims to offer a more tailored investor experience.

In the UK, after 2015 robo-advisors became mainstream with banks and asset management firms, including Barclays, NatWest, and RBS. In 2017, NatWest launched Invest, and its ML algorithm assigns customers a risk category and invests customer money in the appropriate portfolio. In 2020, Barclays launched 'Barclays Plan and Invest', a robo-advisor service developed in partnership with Scalable Capital.²⁵ The service has a minimum £5,000 investment, targets the European market, and offers customers active and passive managed funds through 10,000 possible investment paths.

²⁵ 'Barclays launches robo-advisor for £5K investments with Scalable', Ruby Hinchliffe, *FinTech Futures*, 15 July 2020.

Challenger bank TSB teamed with Wealthify in 2020 to offer investment services based on robo-advising.²⁶

ML techniques have helped financial institutions to increase efficiency in customer interactions and investigate customer buying behaviour. UK Fintech firm FinGenius uses an NLP system that suggests answers as well as interpretations of employee and customer questions. UK-based Fonetic provides speech recognition technology to automate telephone customer interactions.

(v) Algorithmic trading

Since the 1970s, financial institutions have employed algorithms to trade shares and ‘algo-trading’ (AT) is now a dominant force in global financial markets. Also known as ‘automated trading systems’, AT uses complex AI and ML systems to execute extremely fast trading decisions. AT generates 50–70 per cent of equity market trades, 60 per cent of futures trades, and 50 per cent of Treasury market trades (Brummer and Yadav, 2018). To beat competitors, large algorithmic trading is aided by vast amounts of high velocity data. Sentiment analysis can aid ML algorithms in AT by measuring how positive or negative an analyst is about a company’s value.

AT’s benefits include (i) execution of trades at the best possible prices, (ii) increased accuracy and a reduced likelihood of mistakes, (iii) the ability to automatically and simultaneously check multiple market conditions, and (iv) the likely reduction of human errors caused by psychological or emotional conditions (Buchanan, 2019).

Hedge funds, bank proprietary trading desks and houses, and next generation market-makers tend to be algorithmic trading’s target clientele. ML aids AT by including algorithms that include trading decisions, order submission, and managing post-submission orders. Price informativeness should be positively impacted by algorithmic speed (Biais *et al.*, 2011). In 2015, algorithmic trading contributed approximately one-half of UK hedge fund Man Group’s profits.²⁷ The typical hedge fund AT approach is to build volatility, trading costs, asset class, exposure caps, and compliance rules into the algorithm, and this has the advantage of preventing the algorithms breaking the law to make a rushed profit.

Algorithmic trading is responsible for making millions of daily trades. ‘High-frequency trading’ (HFT) is the most recognizable form of AT, as well as its most controversial. Arguments in favour of HFT include a reduction in trading costs, reduced bid–ask spreads, and improved market quality (Hasbrouck and Saar, 2013). For proprietary reasons, financial institutions and hedge funds do not disclose their AI and ML approaches to trading, and ML and DL play an important role in calibrating real time trading decisions. AT usage also involves neural networks, fuzzy logic, and pattern recognition.

AT has its sceptics who are doubtful about the lack of transparency and ‘black box’ nature of AI algorithms. One drawback is that high-frequency traders imitate other traders without exploring the underlying value of the company whose shares are being

²⁶ ‘UK challenger bank TSB partners with Wealthify to launch an investment program’, Omar Faridi, Crowdfund Insider, 4 November 2020.

²⁷ ‘The massive hedge fund betting on AI’, *Bloomberg*, 27 September 2017.

traded (Pasquale, 2015). AT can create instability when trading strategies interact in unforeseen ways, as indicated by the 6 May 2010 ‘flash crash’²⁸ and during the 2016 Brexit referendum, when a flash crash caused the pound sterling to drop by approximately 6 per cent in a 2-minute period.²⁹

Kirilenko and Lo (2013) list the following unintended consequences associated with HFT: flash crashes, fire sales, careless IPOs, cybersecurity breaches, and catastrophic trading errors. They recommend a more systematic and adaptive approach to AT regulation, one in which industrial technological advances are addressed while protecting those who are not. According to a House of Lords 2018 report,³⁰ AI’s full potential can only be realized if potential risks such as algorithmic bias and the opaqueness of black box systems can be mitigated.

(vi) Data and machine learning

In 2006, data was declared the ‘new oil’³¹ and in 2011, the World Economic Forum claimed personal data to be a new asset class. Currently, on average 2.5 quintillion bytes of data are created daily, and it is estimated that 90 per cent of the data available today globally has been created in the past 2 years. By 2025, it is estimated that the data that are created and copied annually (or termed the ‘digital universe’) will be equivalent to 180 zettabytes.

Wall (2017) indicates that a serious limitation to ML processes is data availability for some applications. To address this gap, a rapidly growing trend in ML financial services applications is the role of alternative data (AD). Apart from financial data, alternative data is proving to be an important source of value for many companies. AD extends beyond company filings, earnings calls, or fundamental datasets. It can include (but is not limited to) weather data, satellite tracking data, geo-locational data, app usage, and social sentiment. In other words, alternative data refers to data sources that are not part of the ‘traditional’ set of information metrics used by investment analysts. ML techniques are critical for analysing AD. AD is typically unstructured, very new, and far too large for traditional investment models to interpret.

Whereas traditional data sources may include company financial statements and central bank economic statistics, AD sources extend to include sources as varied as: point-of-sale payment information, satellite tracking and imaging, and social media content. In this sense, AD serves to complement traditional financial and economic data.

Hedge funds adopted AD early on as a means of gaining competitive investment advantage, but current applications of AD extend beyond hedge funds. Traditional

²⁸ The flash crash lasted half an hour during a trading. On 6 May, the Dow Jones Industrial Average 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).

²⁹ ‘UK pound plunges more than 6% in mysterious flash crash’, Jethro Mullen, *CNN Business*, 7 October 2016.

³⁰ ‘Britain urged to take ethical advantage in artificial intelligence’, John Thornhill, *Financial Times*, 15 April 2018.

³¹ This quotation is attributed to Clive Humby.

investment managers, insurance companies, lenders, and other financial service providers are all exploring ways to integrate alternative data into their businesses.

An example of this is Travelex, the UK's largest foreign exchange specialist. In 2020 Travelex suffered significant disruption for approximately a month because it was the target of a cyber-attack.³² Clients' personal identifier information (PII), such as date of birth, national insurance numbers, and credit card information, were stolen. Its big data collaboration network is quite wide. Virgin Money, HSBC, RBS, and Barclays are financial intermediaries who work with Travelex, as well as retailer Asda. Apart from PII and financial information, Travelex also collects sensory and geolocational data. Moneycorp is the second largest foreign exchange broker in the UK. According to its site, Moneycorp uses its data internally to improve products and services and for training purposes. The site is also quite explicit in terms of how it shares its data with Sainsbury Bank.

There are three major data types that can be monetized: raw, anonymized, and synthetic. As of 2020, the use of synthetic data has become more commonplace in the financial services industry. Synthetic data is created when data points are used to generate a new, 'fake' dataset that was not obtained through direct observation. Crucially though, the synthetic dataset still retains the statistical patterns of the real, observed data set. As a result, synthetic data cannot be reverse engineered and, as a synthetic dataset, it does not fall under regulations such as general data protection regulation (GDPR) and the payment card industry data security standard (PCI). Consequently, consumer privacy is protected, and regulations are not violated, meaning that companies can make synthetic data available for research and profit, both within the company and externally. Only synthetic data will ever leave a company's firewall; all the real data remain safe and secure within. UK platform Synthesized is an example of the growing area of synthetic transactional data. It provides synthetic data for testing internal systems, claim prediction performance, bias mitigation, and privacy compliant data.

IV. Machine learning and UK regulation

In the 2019 BoE report, respondents did not believe that regulation would prevent ongoing ML deployment in the financial services industry. Surveyed companies were keen to seek further regulatory guidance regarding ML deployment. In fact, ML-based systems are growing quickly, presenting a challenge to policy-makers. Nevertheless, ML methods hold great promise for improving the credibility of policy evaluation (Athey and Imbens, 2017; Mullainathan and Spiess, 2017). The biggest hurdles appear to be legacy IT systems and data limitations. The 2019 BoE report stresses the need for greater regulatory clarity. Key regulatory themes include data protection and security, accountability, governance, and bias. Regulatory interest in ML is growing to explore explainability, bias, governance, accountability, and data protection.

RegTech is an emerging field that combines big data and ML. Its goals are to make compliance and regulatory activities faster, more efficient, and easier. For example, the

³² <https://www.bbc.co.uk/news/business-51152151>

FCA is looking at the possibility of making its handbook machine-readable and then fully machine-executable. Consequently, this means that ML processes can interpret and implement the rules directly (Citi, 2018).

In 2017, MindBridge and the BoE collaborated to design an AI auditor to help detect transactions and report anomalies. To accomplish these goals,³³ the BoE used ML. ML methodologies can be applied in the context of central banking (Chakraborty and Joseph, 2017). Three ML case studies (two supervised learning cases and one unsupervised learning case) are provided to compare ML methods against traditional econometric methods. Chakraborty and Joseph (2017) examine artificial neural networks, tree-based models, support vector machines, recommender systems, and different clustering techniques. On balance, they find that ML methods generally outperform traditional modelling approaches in prediction tasks. Regarding causal inference properties, this remains an open question.

The GDPR came into force in 2018 and, under this law, EU citizens have the right to receive an explanation for decisions based solely on automated processing, such as those based on ML methods. GDPR stipulates that companies must first obtain consent from an EU citizen before using consumer data. If the EU citizen data is stored on servers located outside of the EU region, GDPR rules apply. Failure to comply with GDPR can result in substantial fines: up to \$22 million or 4 per cent of a company's revenues. In the same year, European MIFID II³⁴ came into effect and requires that firms applying algorithmic models based on AI and ML should have a robust development plan in place. In February 2018 the Prudential Regulatory Authority and the FCA released consultation papers on AT which lists key areas of supervisory focus in relation to MIFID II. During the last 2 years, the UK's Centre for Data Ethics and Innovation (CDEI) has conducted a review into bias in algorithmic decision-making. In July 2020, UK and Australian regulators announced a joint investigation of Clearview AI, the facial recognition company, whose image-scraping tool has been used by police forces around the world.

V. Quantum computing and machine learning

Currently, the UK has 10 supercomputers ranked in the top 500 worldwide and the nation is ranked fifth in Europe behind France, Germany, the Netherlands, and Ireland. In 2020, Nvidia announced it will build the UK's most powerful supercomputer to support AI for healthcare and drug discovery uses.³⁵

The UK is also taking a leading role in quantum computing (QC). QC studies algorithms and systems that apply quantum phenomena to complex problems. In 2019,

³³ 'Bank of England experiments with artificial intelligence and blockchain tech', Brady Dale, *The Observer*, 12 July 2017.

³⁴ Mifid II, or Markets in Financial Instruments Directive, came into effect in 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.

³⁵ 'Nvidia to build UK's most powerful supercomputer', Siddharth Venkatarakrishnan, *Financial Times*, 5 October 2020.

Google announced it had achieved ‘quantum supremacy’, a state where a quantum computer can perform tasks far more rapidly than supercomputers. Traditional computers can only process information in binary format (zeroes or ones). 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 allow the computer to parallel process information. This means that qubits can perform millions of calculations instantly. In the early 1980s, Paul Benioff and Richard Feynman founded the QC field, noting that digital computers (DCs) could not efficiently simulate a probabilistic system, but QC could. QC can potentially process data at speeds far greater than traditional computers.

QCs are much better suited than traditional computers to solving financial problems because QCs operate with random variables, whereas DCs just simulate random variables. The fact that qubits are memory elements that can maintain a linear superposition of both states is a significant change for the financial services industry. Currently, QC is applied to rolling IT security and to portfolio optimization problems ([Rosenberg et al., 2015](#)), with potential applications to finance problems that involve scenario analysis or option pricing ([Lopez de Prado, 2016](#)). 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. Combined with QC capabilities, ML is expected to have a far greater powerful impact ([Lopez de Prado, 2016](#)).

Quantum machine learning has huge potential in the areas of risk profiling, trading optimization, targeting, and prediction. Since 2017, Barclays has been using IBM’s QC software in portfolio optimization applications.³⁶ In September 2020, a £10 million consortium was established to create the UK’s first commercially available quantum computer. The consortium plans to focus on different societal challenges, but UK bank Standard Chartered and the University of Edinburgh are working on quantum machine learning in the finance sector.³⁷ Quantum machine learning is using synthetic data to overcome issues associated with small sample size in some datasets. This boosting of resources means that quantum machine learning will reduce the time to solve complex risk positions. It will also mean that prices on complex financial products determined by quantum machine learning should assist in opening up bigger markets, because an investor will be able to see all the associated risks of a trade. Quantum machine learning also means it would be possible to recalibrate portfolio models more quickly.

VI. What does good AI/ML governance look like?

Good AI/ML governance is one that promotes trust with all key stakeholders and has a framework where risks are properly identified and managed. AI/ML applications and algorithms should be accountable with respect to all desired outcomes. An appropriate governance framework is one with appropriate preventative controls in place and that is regulatory compliant.

³⁶ ‘Can quantum computing transform financial services?’, *The International Banker*, 28 October 2020.

³⁷ ‘Rigetti to build UK’s first commercial quantum computer’, Siddharth Venkataramakrishnan, *Financial Times*, 2 September 2020.

According to the BoE (2019) survey, respondents believe that if ML applications do not work properly, then risk could be amplified if the governance frameworks do not adopt the latest ML technology. Therefore, an effective ML governance framework is one that monitors ongoing performance for the relevant ecosystem, e.g. healthcare, law, finance, etc. It is a framework with continuous fine-tuning and stress-testing to ensure effective risk management. Ultimately, as part of good governance, we need to make sure that we are consistent in understanding, communicating, and prioritizing risks with respect to AI/ML.

Additionally, how do we get more people involved in AI/ML in the context of financial services? First, we need to build on current skills, and support businesses and education. When it comes to finding employees with solid ML skills, a [Statista \(2013\)](#) report states that 86 per cent of those surveyed think ML is the most difficult category in which to place employees. [Hall and Pesenti \(2017\)](#) report on this issue in the UK context. There should be more education that encourages and supports more interdisciplinary approaches in computer science, data literacy, history, and mathematics. There needs to be more diversity of both quantitative and ‘soft skills’. The ethical and social implications of AI/ML should be included within all educational levels. In other words, we need to adapt to a less siloed approach between financial analysts, data engineers, data scientists, risk managers, internal auditors, and ML application developers. In light of the threats to financial stability that develop in a siloed approach, and highlighted by the 2007–8 financial crisis, it is critical that ML not only reduces risk, but that we recognize and implement measures to ensure AI/ML meets society’s needs.

VII. Conclusion—lessons for the future

AI and ML are ‘only scratching the surface’ in terms of potential application to financial services.³⁸ This article discusses the current and emerging applications of these technologies to the financial services industry with a focus on the UK experience. In the near future, two key factors will continue to influence the development of AI and ML: accountability and the effect of Covid-19. Both present opportunities for future research in this area.

In financial services there is an emerging shift away from black box algorithms to ‘white box’ ML algorithms. Neural networks, support vector machines, and tree ensembles are examples of ‘black box’ ML algorithms ([Hoepner et al., 2021](#)). A downside of neural networks is that they are not fully replicable, and lack traceability, transparency, and opportunities to establish causal inference ([Hoepner et al., 2021](#)). Extending the work of [Brooks et al. \(2019\)](#), [Hoepner et al. \(2021\)](#) encourage greater use of ‘white box algorithms’ to encourage explainability and accountability.

The Covid-19 pandemic has forced more retail and financial services firms to adapt to going cashless, mobile payments, and digital currencies. During the Covid-19 pandemic, UK banks like Barclays have seen a spike in robo-advisory service use.³⁹ Within

³⁸ *Financial Times*, 1 July 2020, Nick Huber: AI ‘only scratching the surface’ of potential in financial services.

³⁹ ‘Barclays launches robo-advisor for £5K investments with Scalable’, Ruby Hinchliffe, *FinTech Futures*, 15 July 2020.

the UK, the huge demand for the government business loan schemes (such as CBILS) have forced lenders to provide streamlined digital self-service user journeys.

All financial market participants must consider whether Covid-19 is a tail-risk event or a structural event. If it is a tail-risk event, markets and intermediaries can return to something resembling pre-Covid conditions and expectations. Alternatively, if it is a structural event, financial institutions will need to develop entirely new ML models. At present, the challenge is to avoid ML models becoming biased due to the once-in-a-lifetime Covid-19 event. In the UK, there has been a lack of a systematic framework for gathering track and trace data. In terms of risk management, Covid has contributed to negative ML performance due to underlying data changes (data drift) or the statistical properties of data change (concept drift).

In late 2020, the Bank of England released the results of a survey of the impact of Covid-19 on machine learning and data science. The survey results indicate that around 50 per cent of respondents view ML as being important for future operations because of Covid-19 ([Bholat et al., 2020](#)). Additional survey results indicate that a common use of ML by banks during the pandemic has been to deal with increased customer enquiries. ML also resulted in increased operational efficiency for UK banks in processing a greater volume of government guaranteed loan applications. In terms of planned investment in customer engagement, 60 per cent of UK headquartered banks report that the Covid event has had a positive impact on planned investment in customer applications (with a focus on credit origination and pricing). Around 35 per cent report a negative impact on ML model performance due to shifts in variables such as increasing mortgage forbearance and unemployment.

Appendix I

Timeline of AI and machine learning with financial services applications

- 1842 Ada Lovelace and Charles Babbage propose the creation of computer intelligence.
- 1847 Thornton predicts that computers would one day take over the world.
- 1937 Claude Shannon proposes that Boolean algebra can be used to model electronic circuits.
- 1943 McCulloch & Pitts recognize that Boolean circuits can be used to model brain signals
- 1950 Alan Turing (1950) poses ‘Can Machines Think?’ and 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 Simons create the Logic Machine.
- 1957 [Newell et al. \(1957\)](#) provide the framework to replicate logical flow of human decision making.
- 1958 Frank Rosenblatt introduces a new form of neural network known as ‘perceptron’
- 1961 [Newell and Simons \(1961\)](#) create General Problem Solver
- 1964 Computers understand natural language enough to solve algebraic and word problems

- 1965 Herbert Dreyfus's report severely criticizes 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 ([Lighthill, 1973](#)) ends British government support for AI research
- 1974–80 First 'AI Winter'
- 1980s Early part of decade—Benioff and Feynman create quantum computing
- 1982 In the UK, a new AI report is commissioned and authored by Roger Needham and Peter Swinnerton-Dyer
- 1982 James Simons starts quant investment firm Renaissance Technologies
- 1983–7 The Alvey Programme is the UK's response to the Japanese 5th generation Computer Project
- 1987 Personal Financial Planning System (PFPS) used by Chase Lincoln First Bank; Expert Systems, or Knowledge Systems, emerge as a new field within AI
- 1987–93 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
- 2010 Flash crash occurs on 6 May. In 36 minutes, the S&P crashed 8 per cent, 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
- 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 MiFID II takes effect
- 2018 GDPR takes effect on 25 May
- 2018 UK Department of Work and Pensions tests AI to automate claims processing and combat fraud.

References

- Athey, S., and Imbens, G. W. (2017), 'The State of Applied Econometrics: Causality and Policy Evaluation', *Journal of Economic Perspectives*, **31**(2), 3–32.
- Atiya, A. F. (2001), 'Bankruptcy Prediction for Credit Risk Using Neural Networks: A Survey and New Results', *IEEE Transactions on Neural Networks*, **12**(4), 929–35.

- Bank of England (2019), 'Machine Learning in UK Financial Services', Working Paper, available at: <https://www.bankofengland.co.uk/report/2019/machine-learning-in-uk-financial-services>
- Bauguess, S. W. (2017), 'The Role of Big Data, Machine Learning, and AI in Assessing Risks: A Regulatory Perspective', Keynote Address, OpRisk North America.
- Bellotti, T., and Crook, J. (2009), 'Support Vector Machines for Credit Scoring and Discovery of Significant Features', *Expert Systems with Applications*, **36**(2), 3302–8.
- Berg, T., Burg, V., Gombović, A., and Puri, M. (2020), 'On the Rise of Fintechs: Credit Scoring Using Digital Footprints', *The Review of Financial Studies*, **33**(7), 2845–97.
- Bholat, D. (2020), 'The Impact of Machine Learning and AI on the UK Economy—Conference Overview', Conference on the Impact of Machine Learning and AI on the UK Economy, Bank of England, available at: <https://www.bankofengland.co.uk/events/2020/march/the-impact-of-machine-learning-and-ai-on-the-uk-economy>
- Gharbawai, M., and Thew, O. (2020), 'The Impact of Covid on Machine Learning and Data Science in UK Banking', Bank of England Quarterly Bulletin, available at: <https://www.bankofengland.co.uk/quarterly-bulletin/2020/2020-q4/the-impact-of-covid-on-machine-learning-and-data-science-in-uk-banking>
- Biais, B., Foucault, T., and Moinas, S. (2011), 'Equilibrium High Frequency Trading', in Proceedings from the Fifth Annual Paul Woolley Centre Conference, London School of Economics, May.
- Brooks, C., Hoepner, A. G. F., McMillan, D., Vivian, A., and Wese Simen, C. (2019), 'Financial Data Science: The Birth of a New Financial Research Paradigm Complementing Econometrics?', *The European Journal of Finance*, **25**(17), 1627–36.
- Brummer, C., and Yadav, Y. (2018), 'Fintech and the Innovation Trilemma', *Georgetown Law Journal*, **107**, 235.
- Buchanan, B. (2019), 'Artificial Intelligence in Finance', Alan Turing Institute, London, available at: <https://www.turing.ac.uk/research/publications/artificial-intelligence-finance>
- Cao, C. (2018), 'Quo Vadis? A Comparison of the Fintech Revolution in China and the West', SWIFT Institute Working Paper No. 2017-002, available at <https://swiftinstitute.org/research/quo-vadis-a-comparison-of-the-fintech-revolution-in-china-and-the-west/>
- CB Insights (2019), 'AI Trends to Watch in 2019', available at: <https://www.cbinsights.com/research/ai-trends-2019/>
- Cecchini, M., Aytug, H., Koehler, G. J. and Pathak, P. (2010), 'Detecting Management Fraud in Public Companies', *Management Science*, **56**(7), 1146–60.
- Chakraborty, C., and Joseph, A. (2017), 'Machine Learning at Central Banks', available at: <https://www.bankofengland.co.uk/working-paper/2017/machine-learning-at-central-banks>.
- Citi (2018), 'Bank of the Future: The ABCs of Digital Disruption in Finance', CitiReport, March.
- CognitionX (2018), 'London: The AI Growth Capital of Europe', Mayor of London's Office, available at: https://www.london.gov.uk/sites/default/files/london_theaigrowthcapitalofeurope.pdf
- Cybenko, G. (1989), 'Approximation by Superpositions of a Sigmoidal Function', *Mathematics of Control, Signals and Systems*, **2**(4), 303–14.
- Davis, R. H., Edelman, D. B., and Gammerman, A. J. (1992), 'Machine-learning Algorithms for Credit-card Applications', *IMA Journal of Management Mathematics*, **4**(1), 43–51.
- Desai, V., Crook, J., and Overstreet, G. (1996), 'A Comparison of Neural Networks and Linear Scoring Models in the Credit Union Environment', *European Journal of Operations Research*, **95**(1), 24–37.
- Durand, D. (1941), 'Risk Elements in Consumer Instalment Financing', National Bureau of Economic Research, New York.
- FSB (2017), 'Artificial Intelligence and Machine Learning Financial Services', Financial Stability Board Working Paper, accessed at: <http://www.fsb.org/2017/11/artificial-intelligence-and-machine-learning-in-financial-service/>
- Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2020), 'Predictably Unequal? The Effects of Machine Learning on Credit Markets', *The Effects of Machine Learning on Credit Markets*, 1 October.
- Gammerman, A., and Thatcher, A. R. (1990), 'Bayesian Inference in an Expert System without Assuming Independence', in *Advances in Artificial Intelligence*, Springer.

- 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.
- Hasbrouck, J., and Saar, G. (2013), 'Low-latency Trading', *Journal of Financial Markets*, **16**(4), 646–79.
- Hoepner, A. G. F., McMillan, D., Vivian, A., and Wese Simen, C. (2021), 'Significance, Relevance and Explainability in the Machine Learning Age: An Econometrics and Financial Data Science Perspective', *The European Journal of Finance*, **27**(1–2), 1–7.
- Kaplan, J. (2016), *Artificial Intelligence: What Everyone Needs to Know*, Oxford, Oxford University Press.
- Khandani, A. E., Kim, A. J., and Lo, A. W. (2010), 'Consumer Credit-risk Models via Machine-learning Algorithms', *Journal of Banking & Finance*, **34**(11), 2767–87.
- Kirilenko, A. A., and Lo, A. W. (2013), 'Moore's Law Versus Murphy's Law: Algorithmic Trading and its Discontents', *Journal of Economic Perspectives*, **27**(2), 51–72.
- Liang, D., Tsai, C. F., Dai, A. J., and Eberle, W. (2018), 'A Novel Classifier Ensemble Approach for Financial Distress Prediction', *Knowledge and Information Systems*, **54**(2), 437–62.
- Lighthill, J. (1973), 'Artificial Intelligence: A General Survey', published as part of Artificial Intelligence: A Paper Symposium, Science Research Council, London, SRC, 1–21.
- Lopez de Prado, M. (2016), 'Financial Quantum Computing', available at SSRN: <https://ssrn.com/abstract=2848632> or <http://dx.doi.org/10.2139/ssrn.2848632>
- McCulloch, W. S., and Pitts, W. (1943), 'A Logical Calculus of the Ideas Immanent in Nervous Activity', *Bulletin of Mathematical Biophysics*, **5**, 115–37.
- Makowski, P. (1985), 'Credit Scoring Branches Out', *Credit World*, **74**(2), 30–7.
- Mullainathan, S., and Spiess, J. (2017), 'Machine Learning: An Applied Econometric Approach', *Journal of Economic Perspectives*, **31**(2), 87–106.
- Newell, A., and Simons, H. A. (1961), 'GPS, A Program that Stimulates Human Thought', in H. Billing (ed.), *Lernenende Automaten*, 109–24.
- Shaw, J. C., and Simons, H. A. (1957), 'Empirical Explorations with the Logic Theory Machine', *Proceedings of the Western Joint Computer Conference*, **15**, 218–39.
- Pasquale, F. (2015), *The Black Box Society*, Cambridge, MA, Harvard University Press.
- Proudman, J. (2019), 'Managing Machines: The Governance of Artificial Intelligence', speech at FCA Governance in Banking, available at: <https://www.bankofengland.co.uk/speech/2019/james-proudman-speech-at-fca-conference-on-governance-in-banking-london>
- PwC (2019), 'AI in Financial Services', available at: <https://www.pwc.co.uk/financial-services/assets/pdf/ai-in-financial-services.pdf>
- Rosenberg, G., Haghnegahdar, P., Goddard, P., Carr, P., Wu, K., and Lopez de Prado, M. (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.
- Russell, S. (2021), 'The History and Future of AI', *Oxford Review of Economic Policy*, **37**(3), 509–20.
- Schwab, K. (2017), *The Fourth Industrial Revolution*, Redfern, New South Wales, Currency.
- Sirignano, J., Sadhwani, A., and Giesecke, K. (2021), 'Deep Learning for Mortgage Risk', *Journal of Financial Econometrics*, forthcoming.
- Statista (2013), <https://www.statista.com/search/?q=machine+learning&qKat=search>, New York, Statista.
- (2020), *In-depth: Artificial Intelligence 2020*, New York, Statista.
- Suss, J., and Treitel, H. (2019), 'Predicting Bank Distress in the UK with Machine Learning', Staff Working Paper 831.
- Tang, X., Li, S., Tan, M., and Shi, W. (2020), 'Incorporating Textual and Management Factors into Financial Distress Prediction: A Comparative Study of Machine Learning Methods', *Journal of Forecasting*, **39**(5), 769–87.
- Thomas, L. C. (2000), 'A Survey of Credit and Behavioural Scoring: Forecasting Financial Risk of Lending to Consumers', *International Journal of Forecasting*, **16**(2), 149–72.
- Turing, A. (1950), 'Computing Machinery and Intelligence', *Mind*, **59**(236), 433.

- van Liebergen, B. (2017), 'Machine Learning: A Revolution in Risk Management and Compliance?', *Journal of Financial Transformation*, **45**, 60–7.
- 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>.
- Wang, G., Chen, G., and Chu, Y. (2018), 'A New Random Subspace Method Incorporating Sentiment and Textual Information for Financial Distress Prediction', *Electronic Commerce Research and Applications*, **29**, 30–49.
- Wooldridge, M. (2020), *The Road to Conscious Machines: The Story of AI*, London, Penguin Books.