Legal Aspects of Artificial Intelligence

Richard Kemp

November 2016

KEMP IT LAW

kempitlaw.com

21 Napier Avenue London SW6 3PS Tel 020 3011 1667 Web www.kempitlaw.com

TABLE OF CONTENTS

FIGURES AND TABLE

LEGAL ASPECTS OF ARTIFICIAL INTELLIGENCE[1]

A. INTRODUCTION

1. **Artificial Intelligence in the mainstream**. Writing in the Economist newspaper on 8 October 2016, US President Barack Obama called out artificial intelligence (**AI**) as one of several areas where 'in recent years we have seen incredible technological advances'.[2] Long a backroom area of computer science, AI has captured the popular imagination over the last two years as the range and impact of practical AI applications have expanded at a dizzying pace: a quick search on ft.com for 'artificial intelligence and robotics' returned 4 stories from September and October 2014, 16 for the same period in 2015 and 54 in 2016.

2. **AI and the fourth industrial revolution**. AI is one of several areas of digital innovation that are all both developing increasingly rapidly and interacting with each other in ways whose consequences are challenging to foresee. A useful portmanteau for these changes is the 'fourth industrial revolution'. After steam, electricity and computing, this is the term coined[3] by Davos founder Klaus Schwab for the deep digital transformation now upon us. As digital innovation starts to transform our physical, digital and biological worlds, Mr Schwab's thesis is that we stand on the threshold of vast ranges of IT-driven change where we may expect a 'deep shift' by 2025. Here, AI is just one of a number of technologies which will materially impact all our lives. Others, some of which enable or are enabled by AI, include big data, ubiquitous computing, 'vision as the new interface', implantable and neuro- technologies, the wearable internet, the internet of and for things, connected homes, smart cities and 3D manufacturing.

3. **The 2015 Study Panel's August 2016 'AI and Life in 2030' report**. In September 2016, the AI100 Group (part of the One Hundred Year Study on Artificial Intelligence, a project hosted by Stanford University) published its first report, 'Artificial Intelligence and Life in 2030'.[4] The AI100 Group's remit is to investigate the long-term impact of the science, engineering and deployment of AI-enabled computing systems on people, communities and society, and its core deliverables are five yearly surveys assessing the current state of AI, of which the September 2016 report is the first. The report describes AI and its component parts, reviews AI research trends, overviews AI use cases by sector, and makes recommendations for AI policy.

---

[1] All websites referred to in the footnotes were accessed between 1 October and 5 November 2016.

[2] '*The way ahead*', Barack Obama, The Economist, 8–14 October 2016, page 23 http://www.economist.com/news/briefing/21708216-americas-president-writes-us-about-four-crucial-areas-unfinished-business-economic

[3] '*The Fourth Industrial Revolution*', Klaus Schwab, World Economic Forum, 2016.

[4] https://ai100.stanford.edu/2016-report

4. **What is 'Artificial Intelligence'?** In 1950, Alan Turing, whilst Deputy Director of the Computing Machine Laboratory at the University of Manchester (the developer of the world's first stored-program digital computer) proposed[5] what became known as the Turing Test for calling a machine 'intelligent': a machine could be said to 'think' if a human interlocutor could not tell it apart from another human being in conversation. Six years later, at a workshop to investigate how machines could simulate intelligence at Dartmouth College, New Hampshire, USA, Professor John McCarthy was credited with introducing the term 'artificial intelligence'. A current generally accepted definition is based on two steps, addressing machine intelligence and then the qualities of intelligence:

> "Artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment".[6]

The 'AI and Life in 2030' report (at page 12) also provides a more empirical definition of AI as an operational branch of computer science, stating that 'the field of AI is a continual endeavour to push forward the frontier of machine intelligence' and noting the paradox that when AI ceases to be at the frontier or leading edge it stops being considered AI. More prosaically, 'it's only AI until you know what it does, then it's just software'.

5. **AI research areas**. For many years from the dawn of the computer age, through the mainframe, mini and PC eras and into the early days of the internet, AI research failed to live up to its early promise. Since 2000, however, AI has evolved rapidly and AI-enabled products and services are now becoming established and heralding more widespread commercial success for the future. Driven by exploding volumes of digital data and the advent of the Cloud, major AI developments are taking place in the research areas of machine learning (deep, supervised, unsupervised, reinforcement and large scale machine learning), machine perception (computer vision, speech recognition, natural language processing, Internet of Things) and machine control (autonomous vehicles and robotics). These areas are outlined in Section B of this paper.

6. **The policy perspective**. The 'AI and Life in 2030' report notes (at page 45) that 'AI has the potential to challenge any number of legal assumptions in the short, medium, and long term'. Governments and policy makers around the world are starting to grapple with what AI means for law and policy and the necessary technical and legal frameworks. In April 2016, Japan's Ministry of Economy, Trade and Industry (**METI**) published its Interim Report on the New Industrial Structure Vision as a 'forecasting compass in the public and private sectors to properly address the fourth industrial revolution'.[7] In May, the European Parliament published a draft report on Robotics[8] proposing an advanced robot registration system managed by an EU Agency for Robotics and AI. In the UK, the House of Commons Select Committee on Science and Technology released on 12 October 2016 a

---

[5] '*Computing Machinery and Intelligence*', Alan Turing, Mind, October 1950

[6] '*The Quest for Artificial Intelligence: A History of Ideas and Achievements*', Prof Nils J Nilsson, Cambridge University Press, 2010.

[7] http://www.meti.go.jp/english/press/2016/0427_05.html

[8] '*Draft Report with Recommendations to the Commission on Civil Law Rules on Robotics*', European Parliament Committee on Legal Affairs, 31 May 2016 http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML+COMPARL+PE-582.443+01+DOC+PDF+V0//EN&language=EN

report on Robotics and Artificial Intelligence[9] recommending the establishment of a Standing Commission on Artificial Intelligence and a Leadership Council on Robotics and Autonomous Systems. On the same day, the US Administration published its report on the Future of AI[10] and its AI R&D Strategic Plan.[11]

7. **Scope and aims of this white paper**. This white paper is written from the perspective of the in-house lawyer reviewing the legal aspects of their organisation's first AI project. It:

- addresses in non-technical terms the question: what is AI? and provides a brief outline of current areas of AI research and commercialisation (section B);
- provides three case studies that look at technology, market, legal and regulatory developments in greater depth in each case to give more practical context and perspective on the types of legal issues that arise and how they may be successfully addressed. The case studies are legal services (paragraphs C.15 and C.16) as 'static AI', autonomous vehicles (C.17 and C.18) as 'mobile AI' and smart contracts (C.19 and C.20); and
- reviews at section D the legal aspects of AI from the perspective of regulation (D.23) and agency (D.24), contract (D.25, intellectual property(D.26) and tort law (D.27).

The fourth industrial revolution generally and AI in particular raise profound questions about the nature, timing and extent of the rise of new industries; displacement of employment; societal change; and ethics.

Whilst recognising that political and policy responses to these changes will impact the development of law and regulation as they relate to AI, perhaps in ways that we cannot clearly foresee today, these questions are outside the scope of this white paper.

This paper is written as at 31 October 2016 and from the perspective of English law.

[9] http://www.publications.parliament.uk/pa/cm201617/cmselect/cmsctech/145/14502.htm?utm_source=145&utm_medium=fullbullet&utm_campaign=modulereports

[10] '*Preparing for the Future of Artificial Intelligence*', Executive Office of the President and the National Science and Technology Council (**NSTC**), Committee of Technology, 12 October 2016 https://www.whitehouse.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

[11] '*The National Artificial Intelligence Research and Development Strategic Plan*', NSTC Networking and IT R&D Subcommittee, 12 October 2016 https://www.whitehouse.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf

B. THE ELEMENTS AND TECHNOLOGIES OF AI

8. **Introduction: twinned convergences – the cloud and AI**. Developments in AI have been fuelled by the ability to harness huge tides of digital data. Big data - vast volumes of varied data arriving at velocity[12] - is a product of the Cloud, represented in Figure 1 below as the convergence of data centres and the internet, mobile and social media. Data centres are the engine rooms of the Cloud – where $1bn investments housing 200,000 servers in over a million square feet of space are needed to accommodate current annual growth rates of between 50% and 100% of the four-cloud service provider (**CSP**) majors, Amazon, Microsoft, IBM and Google. Use of the Internet, mobile and social media at scale are in turn driving the Cloud: relative to a global population of 7.4bn in mid-2016, there are currently estimated to be around 6.4bn sensors connected to the internet, 4.6bn mobile users, 3.4bn internet users and 2.3bn social medial users[13]. Rapidly increasing internet, mobile and social media use is in turn fuelling an explosion in digital data volumes, currently growing at a rate of ten times every five years. It is this big data that provides the raw material for AI research and developments.

Figure 1: Twinned convergences: the Cloud and AI

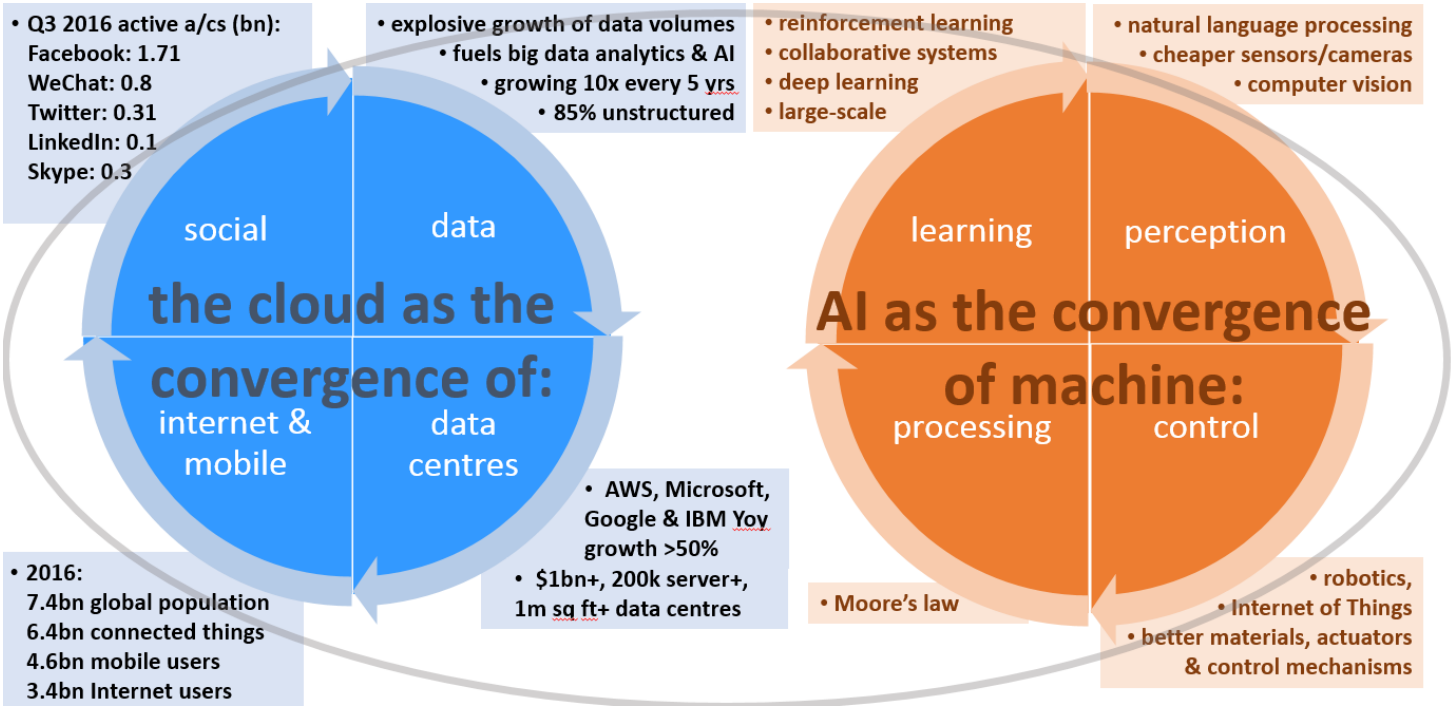

If the Cloud is one side of these twin convergences, AI is the other side. As show in Figure 1, it can be seen as the convergence of four areas of machine capability – processing (paragraph B.9 below),

---

[12] For further information see '*Legal Aspects of Managing Big Data*', Richard Kemp, October 2014 http://www.kempitlaw.com/wp-content/uploads/2014/10/Legal-Aspects-of-Big-Data-White-Paper-v2-1-October-2014.pdf

[13] For the statistics quoted here see http://www.datacenterknowledge.com/special-report-the-worlds-largest-data-centers/ (data centres); https://rcpmag.com/articles/2016/08/02/microsoft-behind-aws-in-cloud.aspx (CSP Q2 2016 market share and growth rates); https://en.wikipedia.org/wiki/World_population (global population); http://www.gartner.com/newsroom/id/3165317 (global internet connected sensors and things);

learning (B.10), perception (B.11) and control (B.12). In the words of Jerry Kaplan in his influential book 'Humans Need Not Apply'[14], what has made AI possible is:

> "the confluence of four advancing technologies … - vast increases in computing power and progress in machine learning techniques … breakthroughs in the field of machine perception … [and] improvements in the industrial design of robots".

9. **Machine processing**. In 1965, Gordon Moore, a co-founder of Intel, famously predicted that the density of transistors (microprocessors) on an integrated circuit (chip) would double approximately every two years. This empirical rule has held good for the last 50 years, meaning in practice is that computer processor speeds have doubled every 18 to 24 months. Although it may be starting run out of steam as processor density starts to produce counter-productive side-effects like excess heat, Moore's law remains the fundamental driver of the computer industry.

10. **Machine learning**. Exponential growth in computer processing power has enabled the techniques of machine learning, by which computers learn by examples and teaching themselves to carry out pattern recognition tasks without being explicitly programmed to do so.

    Deep learning. They do this by using software algorithms to analyse large datasets in what are called neural networks as they seek to mimic the way the human brain works in a process called deep learning. For example, a computer may teach itself to recognise a human face by breaking the inputted data down into layers, where information analysing the problem is passed from layer to layer of increasing abstraction until the final output layer can categorise the entire image as a particular face. Deep learning can be applied to many different applications, domains and use cases and comes in many flavours[15]:

    Supervised learning can be used to train a system with a restricted dataset of labelled examples.

    Unsupervised learning exposes the system to huge volumes of examples and data, but without labelling the data so that the system effectively makes up its own rules as to what to look for and so can discover otherwise hidden correlations.

    Reinforcement learning is a mix of supervised and unsupervised learning, starting with machine learning's training by examples and reduced datasets, and then setting the AI to learn by itself. When Google Deepmind's AlphaGo program beat Lee Sedol, the 18 times Go world champion, in March 2016, AlphaGo's initial machine learning training was then reinforced by playing against itself.

---

https://www.statista.com/statistics/273018/number-of-internet-users-worldwide/ (global internet users); https://www.statista.com/statistics/274774/forecast-of-mobile-phone-users-worldwide/ (global mobile users); and https://www.statista.com/statistics/278414/number-of-worldwide-social-network-users/ (global social media users).

[14] '*Humans Need Not Apply – A Guide to Wealth and Work in the Age of Artificial Intelligence*', Jerry Kaplan, Yale University Press, 2015, pages 38 and 39.

[15] See The Economist Special report: '*Artificial Intelligence – The Return of the Machinery Question*', The Economist, 25 June 2016, page 42 http://www.economist.com/news/special-report/21700761-after-many-false-starts-artificial-intelligence-has-taken-will-it-cause-mass

Large-scale machine learning aims to scale machine learning algorithms to work with much larger datasets so that the algorithms need only operate once (as opposed to several times) on all, or even part of, the data to achieve a faster response.

11. **Machine perception**. Machine learning techniques when combined with increasingly powerful and inexpensive cameras and other sensors are accelerating ***machine perception*** – the ability of processors to analyse data (whether as images, sound, text, unstructured data or any combination) to accurately recognise and describe people, objects and actions.

    Computer vision is currently the most prominent form of machine perception, with applications including face, object and activity recognition and video labelling.

    Speech recognition is another area where machine perception is developing quickly as the error rate has reduced from 25% a few years ago to 6.3% as of September 2016[16] and Amazon, Apple, Google and Microsoft invest heavily in their Alexa, Siri, Google Now and Cortana digital personal assistant systems.

    Natural language processing is emerging as a primary human user interface for AI systems. Enabled by increasing accuracy in voice recognition, systems can respond to one-way user input requests and are starting to interact in two-way conversations. Microsoft's Bing translator enables web pages and larger amounts of text to be translated increasingly accurately in real time and Google in May 2016 announced that 20% of its mobile queries were voice searches.[17]

    Internet of Things (IoT). As the cost of cameras and other sensors continues to decline, increasing numbers of things connect to the internet. In November 2015, research firm Gartner estimated that this number would triple from 6.4bn in 2016 to 20.7bn by 2020[18] as the IoT develops in all its forms – from implantable technologies, the wearable internet and the connected home to autonomous vehicles and Smart Cities. Although these devices generate huge volumes of data, the wide variety of communications protocols and lack of common data standards has generated specific AI research in this area, in the words of the 'AI and Life in 2030' report (at page 13) to 'help tame this tower of Babel'.

12. **Machine control**. Machine control is the design of robots and other automated machines using better, lighter materials and better control mechanisms to enhance the speed and sensitivity of machine response. It adds to the combination of machine learning and machine perception in a static environment the facility of movement in and manipulation of an interactive environment. Essentially, AI is more challenging in a mobile than a static environment and machine control will build on developments in machine learning (particularly reinforcement learning) and perception (particularly force and tactile perception and computer vision). Already by February 2014, 10m units of the

---

[16] '*Microsoft's newest milestone? World's lowest error rate in speech recognition*', Liam Tung, ZDNet.com, 14th September 2016 http://www.zdnet.com/article/microsofts-newest-milestone-worlds-lowest-error-rate-in-speech-recognition/

[17] https://www.linkedin.com/pulse/google-says-20-mobile-queries-voice-searches-nic-shacklock

[18] http://www.gartner.com/newsroom/id/3165317

Roomba autonomous robotic vacuum cleaner had been sold[19] and a survey earlier in 2016 counted more than a dozen commercially available humanoid robots.[20]

13. **Direction of travel of AI development**. The combination of these advances in machine processing, learning, perception and control has set the stage for rapid development of AI. Looking ahead, the 'AI and Life in 2030' report (at page 17) anticipates that significant development will focus on 'human-aware' systems, teaching robots and the IoT:

> "There is a lot of interest in trying to find new, creative ways to develop interactive and scalable ways to teach robots. Also, IoT-type systems - devices and the cloud - are becoming increasingly popular, as is thinking about social and economic dimensions of AI. In the coming years, new perception/object recognition capabilities and robotic platforms that are human-safe will grow, as will data-driven products and their markets."

C. AI SECTOR CASE STUDIES: LEGAL SERVICES, AUTONOMOUS VEHICLES AND SMART CONTRACTS

14. **Introduction: AI in vertical sectors.** Whilst AI can be broken down into its common constituent areas and technologies irrespective of particular use cases, examining the application of AI to particular industry sectors will be helpful to corporate counsel in giving context for reviewing the legal aspects of their organisation's early AI projects. Accordingly, this section works through three case studies:

- AI in legal services as 'static AI' (paragraphs C.15 and C.16);
- autonomous vehicles as mobile AI (C.17 and C.18); and
- smart contracts in the context of blockchain (C.19 and C.20);

highlighting in each case background market and technology developments before considering applicable legal and regulatory aspects.

For other areas, the 'AI and Life in 2030' report (Section II, pages 18 to 41) referred to above includes useful reviews of AI use across eight domains – transportation, home/service robots, healthcare, education, low-resource communities, public safety and security, employment and workplace, and entertainment.

Case Study 1 – AI in Legal Services

15. **AI in legal services: market developments**.

Background: AI and professional services. In their book 'the Future of the Professions', Richard and Daniel Susskind reflect that computerisation of professionals' work began in the 1970s with information retrieval systems and progressed through the 1980s with first generation AI expert systems, the 1990s with knowledge management and the 2000s with Google search. For the 2010s:

---

[19] https://en.wikipedia.org/wiki/Roomba

[20] https://www.smashingrobotics.com/thirteen-advanced-humanoid-robots-for-sale-today/

> "we expect the current decade to be characterised by major progress in Big Data and search. Into the 2020s and beyond, we predict the emergence and adoption of a second wave of AI systems in the professions."[21]

AI adoption in legal services – picking up the pace. Legal services, a £30bn industry in the UK accounting for around 2% of GDP, are representative of the UK professional services world. Driven by cost and competitive pressures, the adoption of machine-learning based AI in legal services has been picking up the pace since mid-2015, 'relieving junior lawyers of time consuming tasks' in the words of a Financial Times article from May 2016.[22] This trend is increasingly visible in B2B legal services, where AI is developing rapidly as a support tool for repetitive, process-intensive, standardisable componentry of higher value legal work across the range of transactional (corporate and finance due diligence, property reports on title), dispute resolution (e-discovery), contract (document assembly, drafting and analysis) and advisory (research and report writing, compliance) service lines. The Table below sets out examples of recently announced B2B legal services AI use cases.

Table - Examples of recently announced B2B AI use cases in the legal services market

Date	Law firm	AI provider	Use case
2015			
Aug	Dentons	IBM/ ROSS Intelligence	Dentons partners with IBM on IBM Cloud. Dentons' NextLaw Labs partners with Ross Intelligence to develop a legal app powered by IBM Watson[23]
Sept	Berwin Leighton Paisner (BLP)	RAVN Systems	RAVN Systems announces that BLP is using its AI platform to manage property Light Obstruction Notices[24]
Oct		Thomson Reuters/ IBM Watson	Thomson Reuters partners with IBM to deliver Watson cognitive computing solutions[25], with Legal as the first use case
Dec	Riverview Law	CIXILEX	Riverview launches the Kim Virtual Assistant built on the CIXILEX platform acquired by Riverview in August.
2016			
May	Baker Hostetler	ROSS Intelligence	Baker Hostetler becomes the first US law firm to license ROSS

[21] '*The Future of the Professions: How Technology Will Transform the Work of Human Experts*', Richard and Daniel Susskind, Oxford University Press, 2015, page 187.

[22] https://www.ft.com/content/19807d3e-1765-11e6-9d98-00386a18e39d

[23] http://www.dentons.com/en/whats-different-about-dentons/connecting-you-to-talented-lawyers-around-the-globe/news/2015/august/dentons-nextlaw-labs-and-ibm-cloud-fuel-legal-tech-startups

[24] https://www.ravn.co.uk/ravn-systems-artificial-intelligence-platform-deployed-successfully-berwin-leighton-paisner/

[25] http://thomsonreuters.com/en/press-releases/2015/october/thomson-reuters-ibm-collaborate-to-deliver-watson-cognitive-computing-technology.html

Date	Law firm	AI provider	Use case
May	BLP	not stated	BLP wins the first contested High Court application to use Predictive Coding in litigation document disclosure[26]
May	Linklaters	RAVN	Linklaters confirms it has signed an MSA with RAVN[27]
June	Allen & Overy	Deloitte	Allen & Overy launch digital derivatives compliance system MarginMatrix with Deloitte[28]
June	DLA Piper	Kira Systems	DLA Piper announces agreement to use Kira in M&A due diligence[29]
July	Clifford Chance	Kira Systems	Clifford Chance announces AI agreement with Kira Systems[30]
Sept	Freshfields	Kira Systems	Freshfields announces agreement to use Kira in its Legal Services Centre[31]
Sept	Slaughter and May	Luminance	Slaughter and May announces collaboration with Luminance on legal due diligence AI[32]

Characteristics of legal sector specific AI applications. What sets this kind of legal services B2B AI apart from more traditional software (like document and workflow automation for example) – what makes it intelligent as well as smart – is its ability to understand, learn and communicate in an unscripted way. Its four main characteristics may be described as:

- ***natural language user interface*** so any user can type or speak their question in plain English;
- ***context aware machine learning*** where the system can be trained on smaller, test ('seed') datasets before being used on much larger datasets for client projects;
- generation of well presented, plain English ***evidence-based responses*** ranked from most to least likely; and
- they are ***cognitive*** – meaning dynamic: it learns from user feedback so next time it is asked the same question its answer will be better.

---

[26] http://www.blplaw.com/expert-legal-insights/articles/blp-wins-contested-application-predictive-coding

[27] https://www.thelawyer.com/issues/online-may-2016/linklaters-becomes-first-magic-circle-firm-to-sign-deal-with-ai-provider-ravn/

[28] http://www.allenovery.com/news/en-gb/articles/Pages/AllenOvery-and-Deloitte-tackle-OTC-derivatives-market-challenge.aspx

[29] https://www.dlapiper.com/en/uk/news/2016/06/dla-piper-partners-with-kira-systems/

[30] https://www.cliffordchance.com/news/news/2016/07/clifford-chance-drives-innovation-strategy-with--artificial-inte.html

[31] http://www.freshfields.com/en/news/Freshfields_continues_innovation_strategy_with_machine_learning_partnership/?LangId=-1

[32] https://www.slaughterandmay.com/news-and-recent-work/news/luminance-launches-with-backing-of-invoke-capital-and-in-collaboration-with-slaughter-and-may/

Virtual assistants and mass market AI developments. Hand in hand with legal sector-specific AI applications go the mass market developments of the global technology companies, especially virtual assistants (Amazon's Echo, Apple's Siri, Google Now and Cortana and Bing from Microsoft). For example, Microsoft's Cortana personal digital assistant on Windows will learn to manage the user's diary, recognise the user's voice, set reminders and answer questions.

16. **AI in legal services: legal aspects**.

***Background - regulatory structure for legal services*.** The regulatory structure for legal services in England and Wales came into effect in October 2011 when most of the Legal Services Act 2007 (**LSA**)[33] came into force. It follows the normal UK pattern of making the provision of certain types of covered services – called 'reserved legal activity' in the LSA – a criminal offence unless the person supplying them is authorised (s.14 LSA). 'Reserved legal activity' is defined at s.12(1) and Schedule 2 LSA and is a shorter list than one might at first imagine[34] so that in fact most 'legal activities'[35] in England and Wales are unregulated.[36] The Legal Services Board (**LSB**) oversees the regulation of lawyers in England and Wales and has appointed eight approved regulators, of which the Solicitors Regulation Authority (**SRA**) is the primary regulator of solicitors[37].

Indirect regulation. In addition to direct legal services regulation, law firms and other legal services providers (**LSPs**) may be indirectly regulated by their client's regulator where that client is itself regulated, for example where it is a bank, by the Financial Conduct Authority (**FCA**) and the Prudential Regulation Authority (**PRA**). This indirect regulation arises through the client regulator's requirements as they apply to the client's contractors, which would include its law firms, and the

---

[33] http://www.legislation.gov.uk/ukpga/2007/29/part/3

[34] 'Reserved legal activities' are (i) rights of audience in the courts, (ii) conduct of litigation in the courts, (iii) preparing documents transferring, or transferring an interest in, land or personal property, (iv) probate activities, (v) notarial activities and (vi) administration of oaths.

[35] Defined at Section 12(3) LSA as covering (i) reserved legal activities and (ii) otherwise in relation to the application of law or resolution of legal disputes, the provision of (a) legal advice and assistance or (b) legal representation.

[36] Contrast the position in the USA for example, where the US State Bar Associations much more zealously protect against the unauthorised practice of law (UPL).

[37] When the LSA came into force, the regulatory functions previously carried out by The Law Society of England and Wales were transferred to the SRA. The Law Society retains its representative functions as the professional association for solicitors. The other LSB approved regulators are (i) the Bar Standards Board (barristers); (ii) CILEx Regulation (legal executives); (iii) the Council for Licensed Conveyancers; (iii) the Intellectual Property Regulation Board (patent and trademark attorneys) as the independent regulatory arm of the Chartered Institute of Patent Agents (CIPA) and the Institute of Trade Mark Attorneys (ITMA); (iv) the Costs Lawyer Standards Board; (v) the Master of the Faculties (notaries); and (vi) the Institute of Chartered Accountants in England and Wales (ICAEW). The ICAEW on 21 July 2016 applied to the LSB to extend the ambit of its approved regulatory functions to cover tax litigation, representation and documents, and notarial and oaths administration services (see https://www.icaew.com/en/about-icaew/news/press-release-archive/2016-press-releases/icaew-applies-to-regulate-further-reserved-legal-services).

In Scotland, solicitors have continued to be regulated by the Law Society of Scotland. The Legal Services (Scotland) Act 2010 in July 2012 introduced alternative providers of legal services as 'licensed legal services providers'. In Northern Ireland, regulatory and representative functions continue to be performed by the Law Society of Northern Ireland.

engagement contract between the client and the law firm, which may flow down contractually certain of the client's regulatory responsibilities and requirements.

The SRA Handbook. The regulatory requirements applicable to SRA authorised LSPs are contained in the SRA Handbook, which sets out the standards and requirements that '[the SRA] expects [its] regulated community to achieve and observe, for the benefit of the clients they serve and in the public interest'.[38] At present, there are no regulatory requirements specifically applicable to AI and the relevant parts of the SRA Handbook are the same Principles[39] and parts of the Code of Conduct[40] that apply in relation to IT systems and services generally.

The ten overarching Principles include acting in the best interests of the client, providing a proper standard of service, complying with regulatory obligations and running 'the business effectively and in accordance with proper governance and financial risk management principles'.

The Code of Conduct is in 15 chapters and sits beneath the Principles setting out outcomes (which are mandatory) and indicative behaviours (which are for guidance). In addition to client care, confidentiality and relationship with the SRA, the relevant outcomes for IT services are mainly at Chapter 7 (management of your business) and include (i) clear and effective governance and reporting (**O(7.1)**), (ii) identifying, monitoring and managing risks to compliance with the Principles (**O(7.3)**), (iii) maintaining systems and controls for monitoring financial stability (**O.(7.4)**), (iv) compliance with data protection and other laws (**O(7.5)**), (v) appropriate training (**O(7.6)**) and (vi) appropriate professional indemnity insurance (**PII**) cover (**O(7.13)**).

SRA Code of Conduct: outsourcing – O(7.10). Specific outcomes are also mandated at **O(7.10)** for outsourcing, which is described in the introduction to Chapter 7 as 'using a third party to provide services that you could provide'. The use of a third party AI platform (but not a platform proprietary to the firm) in substitution for work carried out by staff at the firm is therefore likely to be 'outsourcing' for this purpose. Under O(7.10), a firm must ensure that the outsourcing (i) does not adversely affect compliance, (ii) does not alter its obligations to clients and (iii) is subject to contractual arrangements enabling the SRA or its agent to 'obtain information from, inspect the records … of, or enter the premises of, the third party' provider. This information requirement is likely to be reasonably straightforward to comply with in the case of a third party AI platform used in-house but can give rise to interpretation difficulties for Cloud and other off-premises services.

Client engagement terms - LSPs. As with other IT services, LSPs using AI in client service delivery should consider including express terms around AI use in their client engagement arrangements to set appropriate expectations around service levels and standards consistently

[38] https://www.sra.org.uk/handbook/

[39] https://www.sra.org.uk/solicitors/handbook/handbookprinciples/content.page The Principles have been in effect since October 2011 and were made by the SRA Board under (i) ss. 31, 79 and 80 of the Solicitors Act 1974, (ii) ss. 9 and 9A of the Administration of Justice Act 1985 and (iii) section 83 of the Legal Services Act 2007 with the approval of the Legal Services Board under the Legal Services Act 2007, Sched 4, para 19. They regulate the conduct of solicitors and their employees, registered European lawyers, recognised bodies and their managers and employees, and licensed bodies and their managers and employees.

[40] https://www.sra.org.uk/solicitors/handbook/code/content.page

with SRA duties. SRA regulated LSPs if seeking to limit liability above the minimum[41] must include the limitation in writing and draw it to the client's attention. Firms should therefore consider whether specific liability limitations for AI are to be included in their engagement terms.

Client engagement terms – clients. Equally, clients should insist that their law firms' engagement agreements appropriately document and expressly set out key contract terms around AI services. Clients operating in financial services and other regulated sectors will likely need to go further and ensure that their agreements with the law firms they use include terms that are appropriate and consistent with their own regulatory obligations around (i) security relating to employees, locations, networks, data and records, (ii) audit rights, (iii) continuity, (iv) exit assistance and (v) subcontractors.

PII arrangements. As AI services start to proliferate in the legal services market, it is to be expected that in accepting cover and setting terms and premiums insurers will take a keener interest in how their insured law firms are managing service standard, service continuity and other relevant AI-related risks in client service delivery.

Case Study 2 – Autonomous Vehicles

17. **Autonomous vehicles: technology aspects**.

Vehicle 'autonomy'.[42] If legal services are a prime illustration for the adoption of machine learning-based AI in static services, autonomous vehicles have captured the public imagination as a poster child for AI in motion – the combination of machine learning, machine perception and machine control. 'Autonomous' here means that the vehicle itself is capable with little or no human intervention of making decisions about all its activities: steering, accelerating, braking, lane positioning, routing, complying with traffic signals and general traffic rules and negotiating the environment and other users.

Sensors, digital maps and the central computer. To act autonomously in this way, the vehicle must constantly assess where it is located, the environment and other users around it, and where to move next. These assessments are made and coordinated constantly and in real time by means of sensors, digital maps and a central computer. Figure 2 shows the types of onboard sensors that an autonomous vehicle uses to gather information about its environment, including short, medium and long range radar (**ra**dio **d**etection **a**nd **r**anging), lidar (**li**ght **d**etection **a**nd **r**anging – essentially laser-based radar), sonar (**so**und **n**avigation **a**nd **r**anging), cameras and ultrasound.

In addition to sensors, autonomous vehicles rely on onboard GPS (global positioning system) transceivers and detailed, pre-built digital maps consisting of images of street locations annotated with detailed driving feature information like traffic lights, signs and lane markings. These digital maps are increasingly updated dynamically in real time.

---

[41] LSPs must hold an 'appropriate level' of PII (**O(7.13)**) which under the Insurance Indemnity Rules 2012 must be not less than £3m for Alternative Business Structures ('**ABSs**'), limited liability partnerships ('**LLPs**') and limited companies and £2m in all other cases.

[42] For an excellent guide, see 'Technological Opacity, Predictability, and Self-Driving Cars', Harry Surden (University of Colorado Law School) and Mary-Anne Williams (University of Technology, Sydney), March 2016 - https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2747491

Sense → plan → act. The computer system then receives the data from the sensors, combines it with the map and, using machine learning in a sequential 'sense → plan → act' three step process, constantly (in effect, many thousands of time each second) determines whether, and if so where, when and how, to move. In the ***sensing*** phase, the computer uses the sensors to collect information; in the ***planning*** phase, it creates a digital representation of objects and features based on the data fed by the sensors and aligns the representation to the digital map; and in the ***acting*** phase, the computer moves the vehicle accordingly by activating its driving systems.

Figure 2 – Autonomous vehicles' on board sensors[43]

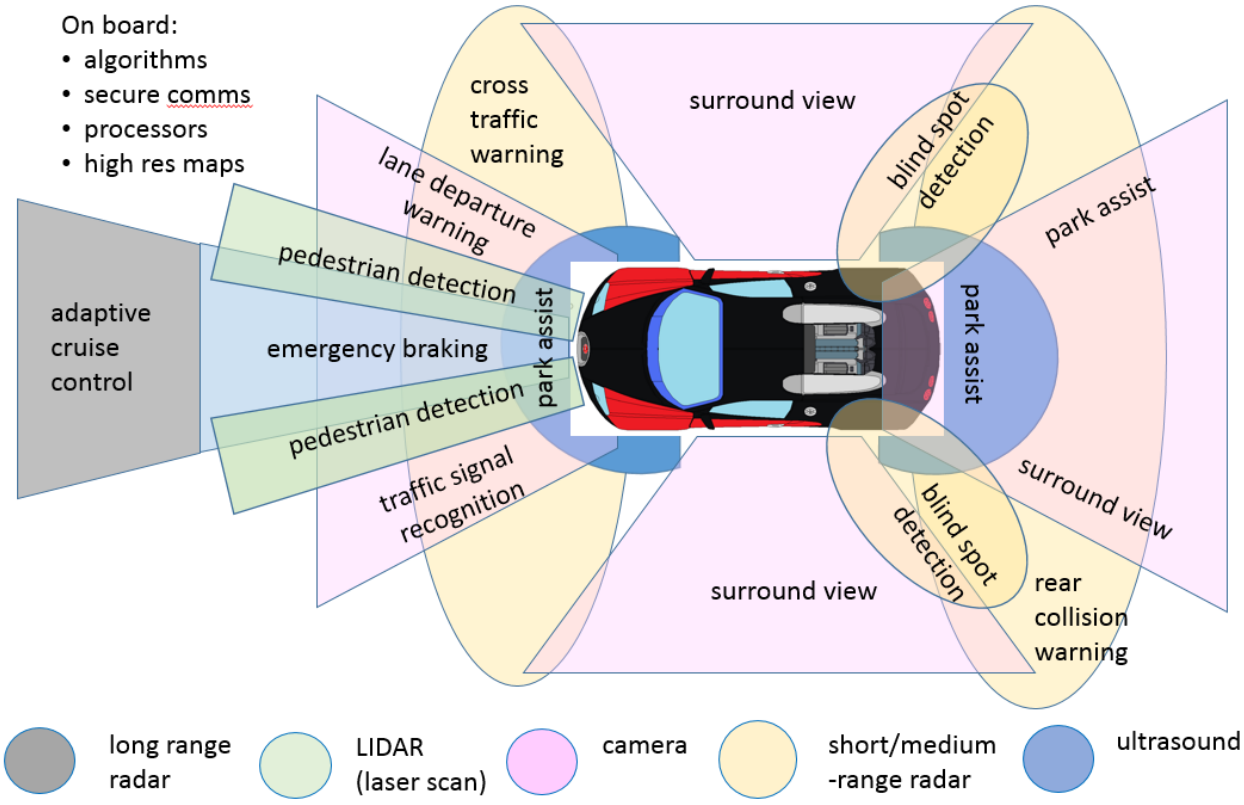

18. **Autonomous vehicles: regulatory aspects**.

The road to autonomous vehicle regulation – issues to be addressed. Since the first of the UK Locomotive ('Red Flag') Acts in 1861, humans have been at the centre of vehicle road driving regulation, whether for speed limits, driving standards, driving licences, vehicle registration or roadworthiness. The removal of human control of motor vehicles that autonomous vehicles predicates is therefore set to transform over 150 years of national and international vehicle, road and traffic legislation and regulation. Key regulatory issues that must be resolved as we move towards road authorisation of autonomous vehicles include (i) connectivity from the vehicle's sensors to other vehicles, objects, road and traffic infrastructure and the environment; (ii) the digital representation of

[43] Source: adapted from OECD International Transport Forum paper, 'Autonomous and Automated Driving – Regulation under uncertainty, page 11, http://www.itf-oecd.org/automated-and-autonomous-driving

the physical world that the vehicle interacts with; (iii) the computer's system for decision making and control; (iv) roadworthiness testing; and (v) all relevant human factors.

SAE International's six levels of driving automation. SAE International (formerly the Society of Automotive Engineers) has mapped[44] out as steps along the road six levels of driving automation, breaking down dynamic driving tasks into the four modes of (i) controlling speed and steering, (ii) monitoring the driving environment, (iii) 'fallback' (failover) performance and (iv) at each level, whether the particular mode of driving is under human or system control.

For the first three levels (no automation, driver assistance and partial automation), the human driver carries out, monitors and is the fallback for the driving modes, with limited automation and system capability for some steering and speed tasks only (like park assist, lane keeping assist and adaptive cruise control).

For the second three levels (conditional, high and full automation) the vehicle progressively takes over steering and speed, driving monitoring, fallback performance, and then some - and finally all - driving modes.

The UK Department for Transport (**DfT**) has conveniently summarised these six levels as moving progressively from (human) '*hands on, eyes on*' through '*hands temporarily off, eyes on*' to '*hands off, eyes off*'.

Pathway to driverless cars - the UK's approach. The DfT has been active in reviewing and preparing for the changes in regulation that will be necessary for autonomous vehicles. It has set up the Centre for Connected and Autonomous Vehicles (**CCAV**) and, under the general heading 'the Pathway to Driverless Cars', has published:

- a detailed review of regulation for automated vehicle technologies[45] (February 2015);
- a Code of Practice for testing[46] (July 2015); and
- a consultation on proposals to support advanced driver assistance systems (**ADAS**) and automated vehicle technology ('**AVT**')[47] (July 2016).

A key challenge for policy makers is that they are aiming at a moving target – regulatory change needs to start now, at a time when it is difficult to predict the future course of development for AVT. The UK is therefore taking a step by step approach:

- first, the February 2015 review confirmed that AVT testing in the UK was already possible: 'you can test automated vehicles on any road in the UK without needing to seek permission from a network operator, report any data to a central authority, or put up a surety bond';

---

[44] SAE International Standard J3016 201401, '*Taxonomy and Definition of Terms Related to on-Road Motor Vehicle Automated Driving Systems*', 16 January 2014, http://standards.sae.org/j3016_201401/

[45] https://www.gov.uk/government/publications/driverless-cars-in-the-uk-a-regulatory-review

[46] https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/446316/pathway-driverless-cars.pdf

[47] https://www.gov.uk/government/consultations/advanced-driver-assistance-systems-and-automated-vehicle-technologies-supporting-their-use-in-the-uk

- second, the testing Code of Practice aims 'to help testers understand how to comply with our laws. It clearly and simply sets out that testers must obey all relevant road traffic laws; test vehicles must be roadworthy; a suitably trained driver or operator must be ready, able, and willing to take control if necessary; and appropriate insurance must be in place';
- third, the DfT is 'actively shaping international vehicle standards to ensure that automated vehicle technologies are safe to come to market';
- fourth, the DfT is starting to tackle the domestic regulatory issues that could inhibit ADAS and AVT uptake. Here, the DfT's July 2016 consultation identified amendments to three related but distinct areas of the UK's current regulatory framework:

  "- **Insurance**: making amendments to primary legislation to ensure insurance products will be available for automated vehicles;

  - **Regulation**: clarifying provisions for the construction and use of near to market technologies (presently, remote control parking) through changes to regulations; and

  - **Highway Code**: providing guidance for drivers about the safe and appropriate use of new ADAS technologies, as well as specific advice in the Highway Code about separation distances for vehicles driving as platoons;"
- fifth, the DfT anticipates future waves of regulatory reform based on lessons learnt 'from real-life experiences of driving of increasingly automated vehicles' and 'from testing fully-automated vehicles both on test tracks and public roads' so 'providing the government with the evidence on which to support future policy decisions'.[48]

Case Study 3 – Smart Contracts

19. **Smart contracts: the blockchain background**.

Background: the blockchain. The blockchain is a comprehensive, always up to date accounting record or ledger of who holds what or who transferred what to whom. The 'what' in the blockchain is virtually anything that can be recorded – physical assets like diamonds and land as well as intangibles like electronic cash (the rationale of Bitcoin, the blockchain's orginator), transactions in securities and other financial instruments, and records of government interaction with citizens.

Blockchain's key features. There are two key features of the blockchain. First, it works through cryptography – authenticating parties' identities and creating immutable hashes (***digests***) of each ledger record, the current page of records (***block***) and the binding that links (***chains***) each block to the earlier ones. This means in theory that the ledger is secure and permanent. Second, instead of one person keeping one instance as 'single version of the truth', the blockchain ledger is distributed: a complete, current copy is held on the computers of each of the network participants (***miner***s) who help keep it up to date. This is deliberate and aims to insulate the integrity of the ledger against cyberattack as any hacker would have to control more than half the network's computing capacity to change any record in the block.

[48] All quotes in this section are taken from the DfT's 'proposals to support ADAS and AVT' July 2016 consultation.

Inhibitors to blockchain uptake. Blockchain is in its infancy, and significant hurdles to commercial adoption remain. First, given the breadth and area of potential applications, many regulatory issues need to be resolved. Second, blockchain is fragmented: there are many different ecosystems[49] that need to agree common standards to all work together. Progress is being made here and in 2016 the ISO has set up Technical Committee 307 on blockchain and electronic distributed ledger technologies.[50] The third hurdle is scalability. The blockchain is enormously power hungry. A report from the National University of Ireland Maynooth in June 2014 famously showed that the total energy consumption of Bitcoin mining was comparable with Ireland's total electricity usage. So more efficient power usage will be key to bigger blockchains.

Blockchain's potential. When these inhibitors have been removed, and in the words of the Economist newspaper, it will "allow strangers to make fiddle-proof records of who owns what".[51] supporters believe that the blockchain's technological integrity will remove the need for the intermediation currently provided by many centralised internet and computer platforms, and enable transactions in a wide range of business areas to be made securely and 'trustlessly' - without any need for any intermediation at all.

Smart contracts. The blockchain paves the way for 'smart contracts', software code representing a self-executing contract as an arrangement that the computer can make, verify, execute and enforce automatically under conditions set in advance. The software can also be used to make and execute chains or bundles of contracts linked to each other, all operating autonomously and automatically. Here, the immutability of the hash (digest) representing each ledger record can get in the way, when all the links in what may end up as a long contractual chain need to execute at the same time to keep the record straight. To get around this, the blockchain is starting to be made editable, with trusted administrators – called oracles – able to change the database.

Benefits of smart contracts. Smart contracts promise a range of benefits including lower costs, latency and error rates (through greater automation, less intermediation and less direct manual involvement) and are likely to enable new business and operating models.

Potential smart contract use cases. Smart contracts represent evolution not revolution. E- and m- commerce today are already predicated on making binding contracts for media, travel and other goods and services through data entry over the internet; and automatic algorithmic trading in financial markets pre-programs AI systems to make binding trades and transactions when certain conditions are satisfied. Smart contracts take this to the next level by further reducing individual human intervention and increasing codification. Areas of potential use include securities and financial instrument clearing and settlement (financial services), insurance claim processing (financial services), electronic patient records (healthcare) and royalty distribution (music and media).

---

[49] The best known are the bank-led R3 (https://r3cev.com/), the exchange-led Post-Trade Distributed Ledger Group (http://www.ptdlgroup.org/)and China's Zhong GuanCun Block Chain Industry Alliance.

[50]http://www.iso.org/iso/home/standards_development/list_of_iso_technical_committees/iso_technical_committee.htm?commid=6266604, International Organisation for Standardisation,

[51] The Economist, 5–11 November 2016, page 10.

20. **Smart contracts: legal and regulatory aspects**. The world of smart contracts can be seen from three perspectives. First, at the developer level, the blockchain smart contract code will need to represent contract law norms; second, the smart contract platform operator will need to contract upstream with the developer and downstream with users; and third, each user will need to contract with the platform operator.

Developer level: Smart contracts in code. In building the smart contract software, the developer will be representing as computer programs a system of normative, contractual rules – a sort of executable 'Chitty on Contracts in code'. The database schema of the system's information architecture – its formal structure and organisation - starts with the flow of information and instructions in the 'real world', takes it through levels of increasing abstraction and then maps it to a data model - the representation of that data and its flow categorised as entities, attributes and interrelationships - in a way that any system conforming to the architecture concerned can recognise and process. Software, as a set of instructions, is not unadjacent to a contract as both set binding rules determining outputs from inputs ('if this, then that'). The information architecture and data modelling of the smart contract system will therefore address, in software, the whole world of contract possibilities that may arise in system use. These include contract formation; payment, performance and lifecycle issues; discharge, liability and resolution; conditionality, dependencies and relief events; audit trail and records generation and retention. The system will also need to cater for relevant regulatory aspects relating to the subject matter of the contracts it is executing – around personal data for example and any consumer and authorisation and compliance regulatory aspects.

Smart contract platform operator level:

- ***Platform operator/developer contract***. At this level, the agreement between the smart contract developer and platform operator is a software 'design, build and operate' agreement – with elements of development, software licensing (if the developer is to retain IP) or transfer (if the IP is to be assigned to the platform operator) and/or service provision with which IT lawyers will be familiar. Particular care will need to be taken in mapping the 'virtual world' of the smart contracts to the real world contractual ecosystem at whose centre sits platform operator. In particular, risk allocation - the 'what if's' of system errors, outages and failures – will need to be managed both contractually (through the governance, service level agreement, liability, indemnity and termination mechanisms) and as appropriate through insurance.
- ***Platform operator/user contract***. The platform operator will need to put in place contract or use terms with each user of the platform. Here, the analogy is with stock exchanges and other trading venues which have detailed membership agreements, contractually binding operational rules, and a range of related agreements and policies regarding software, data licensing and system use and other relevant matters. The platform operator will need to ensure adequate governance and dispute resolution procedures to address the consequences for affected users and counterparties of any failure of the smart contract software anywhere in the ecosystem to operate in the way intended. Regulation of the blockchain and smart contracts is at an embryonic stage, with regulators professing a welcome to the new technology but with little practical progress to date, and operators of smart contract platforms will need to align their systems with these requirements.

User level. The user joining any smart contract system is likely to be presented with a series of more or less standard form contracts that, as a practical matter, it is likely to difficult to change. Key issues for the user include:

- clarity about the extent of contracting authority that the user is conferring on the platform operator's smart contract system – for example, how does it address in all cases where the user is involved the basic issues of contract formation for contracts directly made with the user and any connected agreements on which its own agreements depend;
- evidential requirements (including auditing, record generation/retention and access to/return of data) where the smart contract platform makes and performs commitments entered into in the user's name;
- regulatory issues – control/processing of personal data; system security; regulatory authorisation and compliance requirements - for all/any other platform users, etc; and
- the normal range of contract lifecycle issues, including performance/availability, liability and risk; conditionality/dependencies; and supplier dependence and exit management.

D. LEGAL ASPECTS OF AI

21. **Introduction**. The 'AI and Life in 2030' report quoted above notes (at page 45) that:

> "as a transformative technology, AI has the potential to challenge any number of legal assumptions in the short, medium, and long term. Precisely how law and policy will adapt to advances in AI - and how AI will adapt to values reflected in law and policy - depends on a variety of social, cultural, economic, and other factors, and is likely to vary by jurisdiction."

Just a short list of examples of situations where AI can currently be foreseen gives an idea of the range of legal and regulatory issues that AI will impact:

- a B2B legal services AI applies machine learning incorrectly leading to errors in due diligence;
- a car, ambulance and bus, all operating autonomously, collide at a road intersection;
- separate smart contract systems incorrectly record a negotiated a loan agreement between lender and borrower;
- companies use their AIs in their supply, manufacturing and retail logistics chains;
- construction companies use multiple autonomous machines to build an office block;
- AI is used for the supply of transportation and utilities services in a smart city;
- multiple robots work with each other in the home;
- multiple AI-enabled satellites coordinate with each another in space; and
- medical and health care diagnostics and procedures are planned and carried out by and using AI and robotics.

This section overviews relevant legal and regulatory aspects, aiming to develop an analytical framework that can serve as a checklist of legal areas to be considered for particular AI projects. First, some common misconceptions about AI are clarified (D.22). Regulatory aspects of AI that are set to develop are then outlined (D.23). AI in the areas of agency (D.24), contract (D.25), intellectual property (D.26) and tort (D.27) law are then briefly considered.

22. **Some common misconceptions**. Three misconceptions, all stemming from the fallacy that the embodiment or manifestation of AI has the qualities of a legal person[52], have tended to cloud an analytical approach to the legal aspects of AI, where it is easy to lose sight of the normal tools of legal analysis in the glare of unfamiliar new technology.

    First, the tendency to anthropomorphise AI computer systems (which could be called the 'I Robot fallacy') – to think of robots and AI as analogous to humans and the human brain - can be misleading. A more helpful analogy is with big data and the software that processes that data.

    Second, AI systems, particularly when in motion and especially in popular culture, have been analogised to agents (the 'agency fallacy'), and from there it has only been a short jump to seek to confer rights on and impute duties to these systems as agents. An AI system of itself, at present anyway, is not a legal person and so cannot be an agent.

    As AI 'agents' and smart contracts start to develop and interact, a third misconception (the 'entity fallacy') is to speak of platforms like decentralised autonomous organisations (DAOs) as possessing separate legal personality and able to act independently of their promoters or owners. Generally, as the law currently stands, such an organisation could be expressly incorporated as a separate legal entity as a company or a partnership, where its shareholders or other members would be other legal entities (individuals, companies or trusts, generally). Such a body would behave legally like any other limited liability company or partnership. Failing this, it would likely be classified as a partnership (if carrying on business in common with a view to profit) or an unincorporated association (like a club), in either case in legal terms no different from its members acting together through their partnership or association agreement. This basic distinction is important because, in general terms, a company can, and an unincorporated association cannot, create legal rights and obligations separate and distinct from those of its members. This is not to say that AI will not lead to the development of new types of legal entity – the comparison here would be the development of joint stock companies in the railway age in the UK, when companies were made capable of incorporation first by simple registration and then with limited liability under the Joint Stock Companies Acts 1844, 1855 and 1856.

23. **AI and regulation**. As mentioned in the introduction at paragraph A.6, AI is starting to give governments and policy makers much to grapple with. A number of high level questions arise: first, what interests should AI regulation protect? Should existing regulatory structures be adapted or new ones put in place? How should regulatory burdens be kept proportionate? And what role should central government play? The US October 2016 Future of AI report sets out risk based public protection and economic fairness as the key regulatory interests, using current regulation as the start point where possible and not stifling innovation:

    > "AI has applications in many products, such as cars and aircraft, which are subject to regulation designed to protect the public from harm and ensure fairness in economic competition. How will

[52] The Schedule to the Interpretation Act 1978 defines "person" in the same way as at s.19 Interpretation Act 1889 to "include a body of persons corporate or unincorporated". Persons generally (but not always under English law) have separate legal personality and include individuals (as natural legal persons) and bodies corporate. By s. 1173 Companies Act 2006, "body corporate" and "corporation" "include a body incorporated outside the United Kingdom but do not include (a) a corporation sole, or (b) a partnership that, whether or not a legal person, is not regarded as a body corporate under the law by which it is governed".

> the incorporation of AI into these products affect the relevant regulatory approaches? In general, the approach to regulation of AI-enabled products to protect public safety should be informed by assessment of the aspects of risk that the addition of AI may reduce alongside the aspects of risk that it may increase. If a risk falls within the bounds of an existing regulatory regime, moreover, the policy discussion should start by considering whether the existing regulations already adequately address the risk, or whether they need to be adapted to the addition of AI. Also, where regulatory responses to the addition of AI threaten to increase the cost of compliance, or slow the development or adoption of beneficial innovations, policymakers should consider how those responses could be adjusted to lower costs and barriers to innovation without adversely impacting safety or market fairness."[53]

In the reports from Japan, the EU Parliament, the UK and the US referred to at paragraph A.6, autonomous vehicles are central to the thinking of all these policy makers. The approach of the UK to these the issues is outlined at paragraph C.18 above: first, to ensure there are no obstacles to testing relevant vehicle technology as it develops; then widely consulting with industry and stakeholder groups; and finally to break all elements of necessary regulatory change down into manageable steps (changes to insurance, vehicle construction and the highway code) so as to align regulation pragmatically as ADAS (advanced driver assistance systems) and AVS (automated vehicle technology) develop. This approach is likely to be followed in other countries.

The example of legal services AI at paragraph C.16 above on the other hand shows how use of AI can fit within a current regulatory structure without appreciable change. Here, bought in AI will count as outsourcing under Outcome 7.10 of the SRA Code of Conduct within the UK regime of regulated legal services.

In addition to sector specific regulation, privacy, data protection and cyber security figure large in policy makers' thinking in the context of AI and robotics, largely because it is the combination of big data and machine learning that has been at the root of recent AI developments since 2000. Here, so far as data protection is concerned however, the issues are largely the same as for big data and data protection, and we refer to our white papers on Legal Aspects of Managing Big Data[54] and Big Data and Data Protection[55] for further information.

An important policy question for government is whether to bring together its AI expertise together centrally or whether to decentralise across government departments. The US 'Future of AI' report, in its twenty-three recommendations (pages 40 – 42), appears to advocate common goals for the US federal government and agencies rather than centralisation. In the UK, the October 2016 select committee report notes (at paragraph 72) that "it is too soon to set down sector-wide regulations in this nascent field" and favours the establishment of a standing Commission on AI "with a remit to identify principles to govern the development and application of AI, provide advice to the

---

[53] '*Preparing for the Future of Artificial Intelligence*', Executive Office of the President and the National Science and Technology Council (**NSTC**), Committee of Technology, 12 October 2016, page 11 https://www.whitehouse.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

[54] '*Legal Aspects of Big Data*', Richard Kemp, October 2014, http://www.kempitlaw.com/wp-content/uploads/2014/10/Legal-Aspects-of-Big-Data-White-Paper-v2-1-October-2014.pdf

[55] '*Big Data and Data Protection*', Richard Kemp, November 2014, http://www.kempitlaw.com/wp-content/uploads/2014/10/Big-Data-and-Data-Protection-White-Paper-v1_0-November-2014.pdf

Government, and foster public dialogue" along with a Robotics and Autonomous Systems (**RAS**) Leadership Council.

In the UK, the regulatory position is further complicated as we cross the thresholds of Brexit and the fourth industrial revolution at the same time. Brexit will take up a significant part of the workload of the civil service over the next few years. This will inevitably make it more difficult at this seminal time for AI policy makers to take effective and high quality decisions. The UK IT industry will need to shout loud on AI issues to make its voice heard above the Brexit din.

24. **AI and agency law**. Agency is a relationship between two legal persons. In the words of the leading work on UK agency law, it is:

> "the fiduciary relationship which exists between two persons, one of whom expressly or impliedly manifests assent that the other should act on his behalf so as to affect his relations with third parties (the principal) and the other of whom similarly manifests assent so to act or so acts pursuant to the manifestation."[56]

As mentioned at paragraph D.22, a common misconception is to regard AI systems as 'agents' and so capable of acting for their 'principal'. An AI-enabled or autonomous system is not of itself a legal person. It – or rather, the personal property (as goods) or intellectual property rights (as software, etc) in relation to it – belongs, or is provided as a licence or as a service, to its owner or the person using it. An AI system could be operated as limited liability company or other legal entity, but then it is that entity that owns it and licenses it or otherwise contracts for its use.

25. **AI and contract law**. Commercial contracts for the development, provision and use of B2B AI systems between developer/licensor/provider and licensee/customer will, in the short term, be largely indistinguishable from other software contracts, whether provided on-premise as a licence or in the Cloud as a service. Similar issues to those in software agreements will need to be addressed in AI agreement and are not considered further here.[57] Equally, mass market B2C AI services (like Amazon's Alexa, Apple's Siri, Google Now and Microsoft's Cortana digital personal assistants) will continue to be made available to subscribers on the basis of click accept licensing terms.

The legal analysis starts to get more complex in the case of smart contracts. Paragraph C.20 above overviews contractual aspects from the standpoint of the developer, smart contract platform operator and user. Blockchain enabled smart contracts will have the ability to make virtually real time interlocking chains of contracts linked by dependencies. For each link in the chain the basics of contract formation in the jurisdiction(s) that are agreed to govern the smart contract ecosystem will need to be present, both as code and contract - the software code that implements the system and in the agreement(s) governing use. In the UK these include (i) that each contracting party has the necessary legal capacity; (ii) intention to create legal relations; (iii) offer; (iv) acceptance; (v)

---

[56] '*Bowstead and Reynolds on Agency*', 20th Edition, page 1, Sweet & Maxwell, 2016

[57] For further information on IT contracts generally see our white paper '*Demystifying IT Law*', 7 March 2016 http://www.kempitlaw.com/wp-content/uploads/2016/02/Demystifying_IT-White-Paper-KITL-v1.2-March-2016.pdf and in relation to Cloud legal services, our blog '*Law Firms and Contracting for the Cloud*', 15 October 2014, http://www.kempitlaw.com/law-firms-and-contracting-for-the-cloud/

communication of acceptance; (vi) consideration; (vii) obligations recognised by law; and (viii) certainty of terms.[58]

Where the chain of contracts becomes extended, the possibility arises that an earlier contractual link will be broken, for example, because the contract formation requirements were not met or the contract was discharged through breach. The impact of a broken upstream contractual link on a downstream contract in an AI-enabled or smart contract system is likely to raise novel contract law questions. An agreement may lack contractual force for uncertainty[59] or any downstream contractual link in the chain may be dependent – as a condition precedent – on the performance of all logically prior, upstream agreements. An almost limitless range of possibilities will need to be addressed in software terms in the smart contract code base and covered in the express contractual terms of the 'house rules' that govern use of the system. It is therefore foreseeable that we will see contract law evolve in this area over the next few years as novel disputes thrown up by smart contract systems are settled through the courts.

26. **AI and intellectual property**. AI will provide a significant impulse to the development of intellectual property law, particularly as machine and cognitive learning start to enable computers to generate new works and invent novel ways of doing things.

In the field of copyright, UK law has always developed with new bits added on Lego-like as technology evolves.[60] A key question that arises concerns ownership of copyright works generated by AI systems, where s.9(3) of the UK Copyright Designs and Patents Act 1988[61] (**CDPA**) provides that:

> "In the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken"

and 'computer-generated' is defined at s.178 as meaning:

> "that the work is generated by computer in circumstances such that there is no human author of the work."

The operative terms here are fraught with difficulty and, in the absence of significant case law on the point to date to clarify for example what is meant by 'undertaking necessary arrangements' for the creation of the work where 'there is no human author', widespread use of AI systems is likely to lead to legal developments in this area. Accordingly, parties to agreements for the development and use of an AI system that may be expected to result in new copyright works should consider including any necessary express terms as to their ownership, assignment and licensing.

---

[58] See *Chitty on Contracts*, 32nd Edition, Parts One and Two, Sweet & Maxwell, 2015.

[59] See Chitty, paragraph 2-147

[60] Software was first given literary copyright protection in 1985 in the UK by the Copyright (Computer Software) Amendment Act 1985 and copyright aspects of the internet were introduced into English law by the Copyright and Related Rights Regulations 2003 (SI 2003/2498), implementing the EU Directive 2001/29/EC on Copyright and Related Rights in the Information Society.

[61] http://www.legislation.gov.uk/ukpga/1988/48/contents

Equally, use of AI systems may result in new inventions and the question arises whether these computer implemented inventions are capable of patent protection. S.1(2)(c) Patents Act 1977 (**PA**)[62] excludes "a program for a computer" from patent protection to the extent that the patent application "relates to that thing as such". This has led to a line of cases in the UK over the last ten years or so[63] which has sought to establish and clarify a test for determining the contribution that the invention makes to the technical field of knowledge (potentially patentable) beyond the computer program "as such" (not patentable).

If the invention is potentially patentable on this basis, s.7(3) PA provides that:

> "In this Act "inventor" in relation to an invention means the actual deviser of the invention and "joint inventor" shall be construed accordingly"

and s.7(2)(a) provides that a patent for invention may be granted "primarily to the inventor or joint inventors". US law is more specific in defining (at 35 USC §100(f) and (g)) "inventor" as "the individual or, if a joint invention, the individuals collectively who invented the subject matter of the invention". The context of s.7(3) means that the 'actual deviser of the invention' should be a 'person' and there is no regime similar to that for copyright for computer-generated works. Again, the take away from the patent law perspective is that it is worth legislating in B2B AI contracts for ownership, assignment and licensing aspects of AI generated inventions and patent rights as well as copyright works.

27. **AI and tort law**. Outside regulatory and statute law, it is perhaps the common law area of tort that is likely to see the most important AI-influenced legal developments. Product liability will be evidently also be relevant for autonomous vehicles, robots and other 'mobile' AI-enabled or autonomous systems, and the tort of breach of statutory duty will also be relevant depending on the regulatory backdrop.

    'Static' and 'mobile' AI are likely to involve their providers and users in common law duties of care (negligence) and nuisance and escape (*Rylands* v *Fletcher*) liability. Negligence under English law centres on the existence of a duty at common law 'to be careful'. The list of situations giving rise to a duty of care is not fixed: in the famous words of the UK House of Lords "the categories of negligence are never closed"[64], and it is hard to imagine that the common law duty of care will not arise in relation to many, or most, kinds of AI.

    Nuisance and escape (*Rylands* v *Fletcher*) liability are based on interference with the use or enjoyment of land, and are more likely to be relevant for robots, autonomous vehicles and other kinds of 'mobile AI' than for 'static AI' systems. If a robot runs amok, the situation may be analogised to straying animals where under English law liability has been codified by statute under the Animals Act 1971, s.4 of which for example imposes strict liability for straying animals. This points back to statutory regulation of AI but, for the moment, one can easily imagine the common law being

[62] https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/354942/patentsact1977011014.pdf

[63] Starting with *Aerotel Ltd* v *Telco Holdings Ltd* and *Macrossan's Patent Application* [2006] EWCA Civ 1371

[64] Lord Macmillan in *Donoghue* v *Stevenson* [1932] A.C. 562 at p. 619.

extended to treat AIs causing unreasonable annoyance to a neighbours as nuisance in the same way as for animals.

The rule in *Rylands* v *Fletcher* is that:

> "a person who for his own purposes brings on his lands and collects or keeps there anything likely to do mischief it if escapes must keep it in at his peril, and if he does not do so, is prima facie answerable for all damage which is the natural consequence of its escape."[65]

The principle extends 'dangerous things' as 'things' 'likely to do mischief' on escape and has been applied to motor vehicles and electricity but not an aeroplane or a cricket ball driven out of the ground[66]. Extending *Rylands* v *Fletcher* escape liability in tort to AI would therefore appear to be a simple extension consistent with past decisions.

E. CONCLUSION

28. **Conclusion**. As we approach 2017, AI is a central part of the deep digital transformation of the fourth industrial revolution whose threshold we are now crossing. As big data fuelled machine learning, perception and control develops, AI will come to impact our family, home and working lives perhaps as much as any industrial change over the last 250 years. Early examples of AI in action include AI in legal services, autonomous vehicles and smart contracts. AI will challenge legal assumptions short, medium and long term. Policy makers and regulators are consequently starting to grapple with what AI means for law and policy and the necessary technical and legal frameworks. 2016 has seen important policy announcements in the EU, Japan, the UK and the USA. In order to successfully manage AI projects, lawyers in the field will need to keep up to date with AI related regulatory and policy developments, and developments in contract, intellectual property and tort law as the courts decide disputes and make new case law. AI promises to be another fascinating area for IT lawyers.

Richard Kemp,
Kemp IT Law,
London,
November 2016
richard.kemp@kempitlaw.com
Tel: 020 3011 1670

[65] (1866) L.R. 1 Ex. 265 at 279.

[66] Motor car – *Musgrove* v *Pandelis* [1919] 2 K.B.43; electricity – *National Telephone Co.* v *Baker* [1893] 2 Ch. 186; aeroplane – *Fosbrooke-Hobbs* v *Airwork Ltd* [1937] 1 All E.R. 108; cricket ball – *Bolton* v *Stone* [1951] A.C. 850.