Demystifying Artificial Intelligence

Richard Kemp

June 2017

IT law & regulation,
4thIR, Brexit & IT law,
Cloud computing law, Data law,
IT&FS, Software law, standards

blog, vlog, alerts,
white papers
and events

KEMP IT LAW

IT Law at the Apex

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

DEMYSTIFYING ARTIFICIAL INTELLIGENCE

TABLE OF CONTENTS

TABLE OF FIGURES

DEMYSTIFYING ARTIFICIAL INTELLIGENCE

A. INTRODUCTION: 'IT'S ONLY AI UNTIL YOU KNOW WHAT IT DOES, THEN IT'S JUST SOFTWARE'

1. **Significance of AI.** In research consultancy Gartner's '*Top 10 Strategic Technology Trends for 2007*' survey,[1] Gartner Vice-President and Fellow David Cearley said

> "**over the next 10 years, virtually every app, application and service will incorporate some level of AI. This will form a long-term trend that will continually evolve and expand the application of AI and machine learning for apps and services.**"

AI therefore looks set to become a big part of in-house legal workloads. To combat the inevitable FUD factor, a good way of getting to grips with it and making AI more accessible is to remember that 'it's only AI until you know what it does, then it's just software'. This note aims to continue the demystification process.

2. **Twinned convergences: the Cloud and AI.** Fuelled by exploding volumes of big data – digital data is growing at a compound rate of 60% per year – AI can be seen (as shown in Figure 1) as one side of twinned convergences, whether the other side is the Cloud and its constituent parts (data centres – the Cloud's engine room – social, internet, mobile and data).

Figure 1 – Twinned Convergences: The Cloud and AI

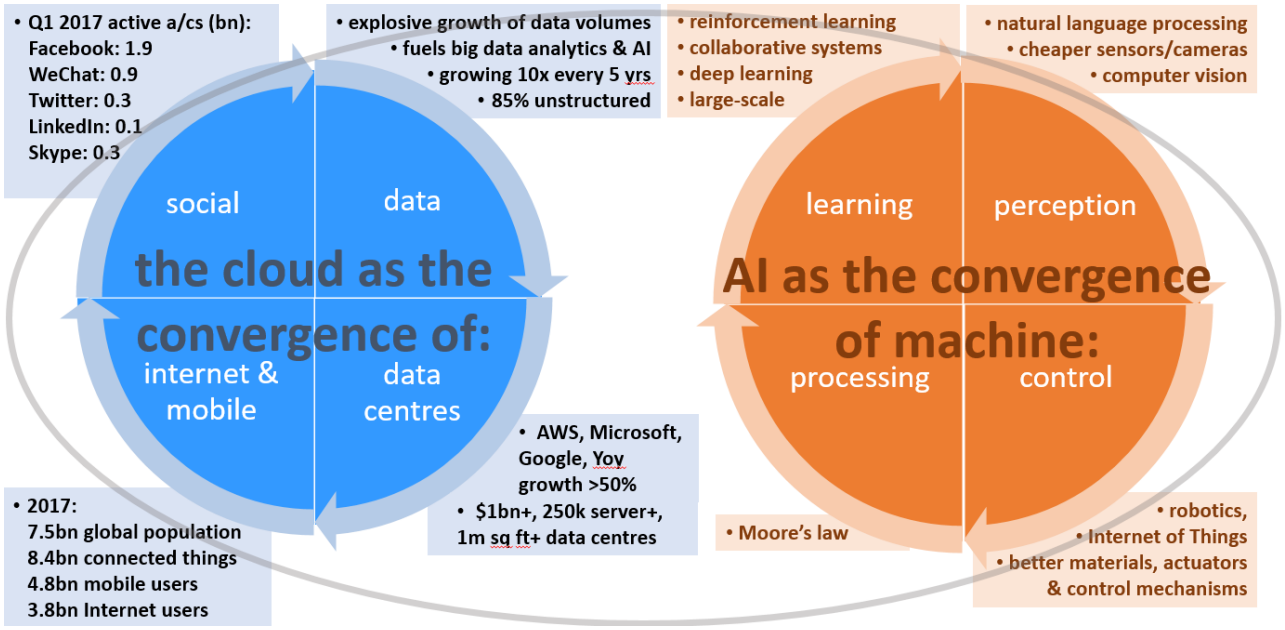

[1] '*Gartner Identifies the Top 10 Strategic Technology Trends for 2017*', 18 October 2016 http://www.gartner.com/newsroom/id/3482617

B. AI AS THE CONVERGENCE OF MACHINE PROCESSING, LEARNING, PERCEPTION AND CONTROL

3. **AI as machine processing, learning, perception and control**. AI can be seen from this perspective as the convergence of machine processing, learning, perception and control. Exponential growth in ***machine processing*** (Moore's law, that processor power doubles every two years) power has enabled the techniques of ***machine learning***, by which computers learn by examples and teach themselves to carry out pattern recognition tasks without being explicitly programmed to do so. Combining machine learning with billions of Internet–connected sensors enables ***machine perception*** – advances in implantable and wearable devices, personal digital assistants, the Internet of Things ('**IoT**'), connected homes and smart cities. Add actuation – the ability to navigate the environment – to static machine learning and perception and you get ***machine control*** – autonomous vehicles, domestic robots and drones.

4. **Machine learning**. Deep learning, a machine learning technique, is emerging as AI's 'killer app' enabler. It works by first using large training datasets to teach AI software to accurately recognise patterns from images, sounds and other input data and then, once trained, the software's decreasing error rate enables it to make increasingly accurate predictions. Deep learning is the core technology behind the current rapid uptake of AI in a wide variety of business sectors from due diligence and e-discovery by law firms to the evolution of autonomous vehicles. To show how this happens, Microsoft in October 2016 released an updated version of Cognitive Toolkit, its deep learning acceleration software, and provided in its accompanying[2] blog an example (reproduced below) of how the toolkit used training sets to increase training speech recognition accuracy.

Figure 2 – Microsoft Cognitive Toolkit:
increasing speech recognition accuracy by epochs of training set use

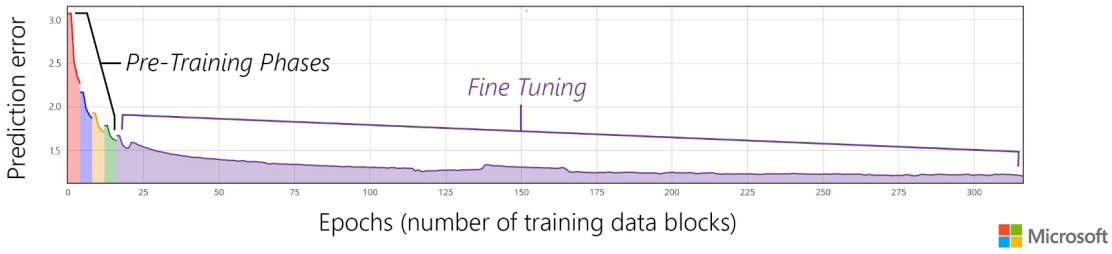

This pattern – using the machine learning software to reduce prediction error through training and fine tuning, then letting the software loose on the workloads it's to process – is at the core of AI in professional services. It's what's behind the AI arms race in law (standardising componentry of higher value work like due diligence, e-discovery in litigation, property reports on title, regulatory

[2] '*Microsoft releases beta of Microsoft Cognitive Toolkit for deep learning advances*', 25 October 2016 http://blogs.microsoft.com/next/2016/10/25/microsoft-releases-beta-microsoft-cognitive-toolkit-deep-learning-advances/#sm.0000lt0pxmj5dey2sue1f5pvp13wh

compliance), accountancy (audit processes, tax compliance, risk) and (coupled with IoT sensors) insurance, for example.

5. **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[3] 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.[4]

   - ***Internet of Things (IoT)***. As the cost of cameras and other sensors continues to decline, increasing numbers of things connect to the internet. Estimates vary, but it's quite possible that there are around 10 billing connected things at the moment, and that this will rise to 50 billion by 2020 as the IoT develops in all its forms – from implantable technologies, the wearable internet and the connected home to autonomous vehicles and Smart Cities.

6. **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, a type of machine learning) and perception (particularly force and tactile perception and computer vision). Already by February

---

[3] '*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/

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

2014, 10m units of the Roomba autonomous robotic vacuum cleaner had been sold[5] and a survey from 2016 counted more than a dozen commercially available humanoid robots.[6]

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

This section briefly overviews AI at work in three sectors: legal services, autonomous vehicles and smart contracts in the context of blockchain.

7. **AI in legal services.** The law, a £30bn industry in the UK accounting for around 2% of GDP, is representative of the UK professional services world. Driven by cost pressures and regulatory liberalisation, the adoption of machine-learning based AI in legal services has been speeding up since mid-2015, "relieving junior lawyers of time consuming tasks" in the words of a Financial Times article from May 2016.[7] 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. What sets this kind of legal B2B AI apart from more traditional software 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 plain English ***ranked evidence-based responses***, from most to least likely; and
- it is ***cognitive*** – meaning dynamic: it learns from user feedback so next time it is asked the same question its answer will be better.

The UK legal services regulatory regime is liberal for AI. Essentially for regulated firms, 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 likely to be treated as 'outsourcing'. Under the SRA Code of Conduct[8], Outcome 7.10, a firm must ensure that the outsourcing:

- does not adversely affect compliance;

---

[5] https://en.wikipedia.org/wiki/Roomba
[6] https://www.smashingrobotics.com/thirteen-advanced-humanoid-robots-for-sale-today/
[7] https://www.ft.com/content/19807d3e-1765-11e6-9d98-00386a18e39d
[8] https://www.sra.org.uk/solicitors/handbook/code/content.page

- does not alter its obligations to clients; and

- 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 requirement will be straightforward to comply with for AI used in-house but can give rise to difficulties in the Cloud.

At the moment, using AI in legal services isn't easy: getting the right input data, getting the dataset tuning right, and then correctly applying the trained dataset to the production workload - none of this is plain sailing. Contractually, in terms of the statement of work between firm and client, the accent should be on collaboration and setting realistic expectations and outcomes. Currently, much AI work is effectively in 'beta', and this will be reflected from the firm's point of view in lower service (and likely lower fee) levels, with the expectation that as more experience is gained so service levels and fees will rise. Firms' and users' professional indemnity insurers are getting more interested in AI risks in professional services, so an early conversation to let them know and get advance notice of anything they particularly care about may not be out of place.

8. **Autonomous vehicles[9].** 'Autonomous' here means that the vehicle itself is capable with no human intervention of making decisions about all its activities: steering, accelerating, braking, lane positioning, routing, complying with traffic signals and traffic rules, and negotiating the environment and other users. 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, in real time by sensors, digital maps and central computers.

   Figure 3 below 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, the really clever bit of autonomous vehicles), sonar (**so**und **n**avigation **a**nd **r**anging), cameras and ultrasound.

   In addition to sensors, autonomous vehicles rely on onboard GPS (global positioning systems) 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.

   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

[9] 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

aligns the representation to the digital map; and in the ***acting*** phase, the computer moves the vehicle accordingly by activating its driving systems.

The removal of human control of motor vehicles that autonomous vehicles predicates will transform vehicle road driving regulation, which has grown up around speed limits, driving standards, driving licences, vehicle registration and roadworthiness. Key regulatory issues that must be resolved as we move towards road authorisation of autonomous vehicles include:

- conn

SAE International (formerly the Society of Automotive Engineers) has mapped[11] out as steps along the way six levels of driving automation, breaking down dynamic driving tasks into the four modes of:

- controlling speed and steering;
- monitoring the driving environment;
- 'fallback' (failover) performance; and
- at each level, whether the particular mode of driving is under human or system control.

The UK Department for Transport (**DfT**) has been active in reviewing and preparing for the changes in regulation that will be necessary for autonomous vehicles and has conveniently summarised the six levels of driving automation as moving progressively from (human) '*hands on, eyes on*' through '*hands temporarily off, eyes on*' to '*hands off, eyes off*'.

9. **Smart contracts and the blockchain.** The background to smart contracts is the blockchain, as 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 originator), transactions in securities and other financial instruments, and records of government interaction with citizens. Blockchain has two key features. 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 (miners) 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,

In the words of the Economist newspaper, the development of the blockchain will "allow strangers to make fiddle-proof records of who owns what".[12] 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.

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

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

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.

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. 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).

The world of smart contracts can be seen from three perspectives:

- at the ***developer level***, the blockchain smart contract code will need to represent contract law norms - a sort of executable 'Chitty on Contracts in code';

- at the ***platform developer/platform operator*** contract level, the agreement 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; and

- at the ***platform operator/user contract*** level, 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.

D. LEGAL ASPECTS OF AI

10. **Introduction.** Increasingly rapid adoption of AI over the next five years will bring challenges for law and policy makers as the law struggles to keep up. A couple of initial 'do's' and 'don'ts':

- don't anthropomorphise AI: in legal terms, AI is personal property not a person (what you might call the '**I Robot fallacy**'), AI systems aren't 'agents' in any legal sense (the '**agency fallacy**') and AI platforms themselves don't possess separate legal personality (the '**entity fallacy**').

- don't be blinded by the glare of the new and do go back to first principles – whether it's regulation or in contract, tort or copyright law.

11. **Data law.** Data law is now right at the centre of AI. In the run up to May 2018 when both the General Data Protection Regulation and the Network and Information System Directive take effect, data protection and data security are rising up the business agenda as firms prepare themselves for a new data-centric world. But it's not just privacy and security – legal rights and duties around data

licensing (have I got the right permissions to do what I'm doing with my data?) and data sovereignty (my data in your data centre) are also becoming more important.

12. **Regulation.** Regulators around the world are grappling with how to address AI. What happens when an autonomous car and bus collide? Or when smart contract systems incorrectly record a negotiated mortgage or personal loan agreement? Or when AI-enabled due diligence misses the point? The emerging consensus on approach involves a number steps: establishing governmental advisory centres of AI excellence; adapting existing regulatory frameworks to cater for AI where possible; and (perhaps) some system of registration for particular types of AI.

13. **Intellectual property law.** 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, a key question that arises concerns ownership of copyright works generated by AI systems, where the way in which UK copyright law[13] deals with 'computer-generated' works is fraught with difficulty and is likely to lead to litigation. Accordingly, parties to agreements for the development and use of an AI system that may be expected to result in new copyright works should include appropriate express terms on ownership, assignment and licensing. Equally, use of AI systems may result in new inventions, such computer-implemented inventions may be capable of patent protection and contractual parties should again expressly provide in their AI contracts for ownership, assignment and licensing of AI generated inventions and patent rights.

14. **Tort law.** Tort law is the area 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*[14]) 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"[15], 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.

15. **Liability theories.** We're likely to see rapid and complex developments around legal theories of liability as AI becomes the norm. For those involved in the value chain for autonomous vehicles for

---

[13] http://www.legislation.gov.uk/ukpga/1988/48/contents
[14] (1866) L.R. 1 Ex. 265 at 279.
[15] Lord Macmillan in *Donoghue* v *Stevenson* [1932] A.C. 562 at p. 619.

example, these issues are likely to be acute: accidents caused by driverless cars are likely to provoke sharp public outcry, and managing all relevant issues (from deployment of algorithms and software, ethical considerations, evidencing the choices made at the time, liability theories, risk management, insurance and compensation to managing public opinion) will figure large in business planning over the next few years.

What is the appropriate balance of rights and responsibilities between an AI user and its AI software provider or cloud service provider? What is negligence (breach of the common law duty of care) in AI terms? And how will the state, and industry regulators intervene to manage AI and AI liability? We can expect answers to these questions in the months and year ahead.

Richard Kemp
Kemp IT Law, London
richard.kemp@kempitlaw.com
June 2017