CMS
law·tax·future

AI in Life Sciences

Legal perspectives on the opportunities and challenges of AI for life sciences companies

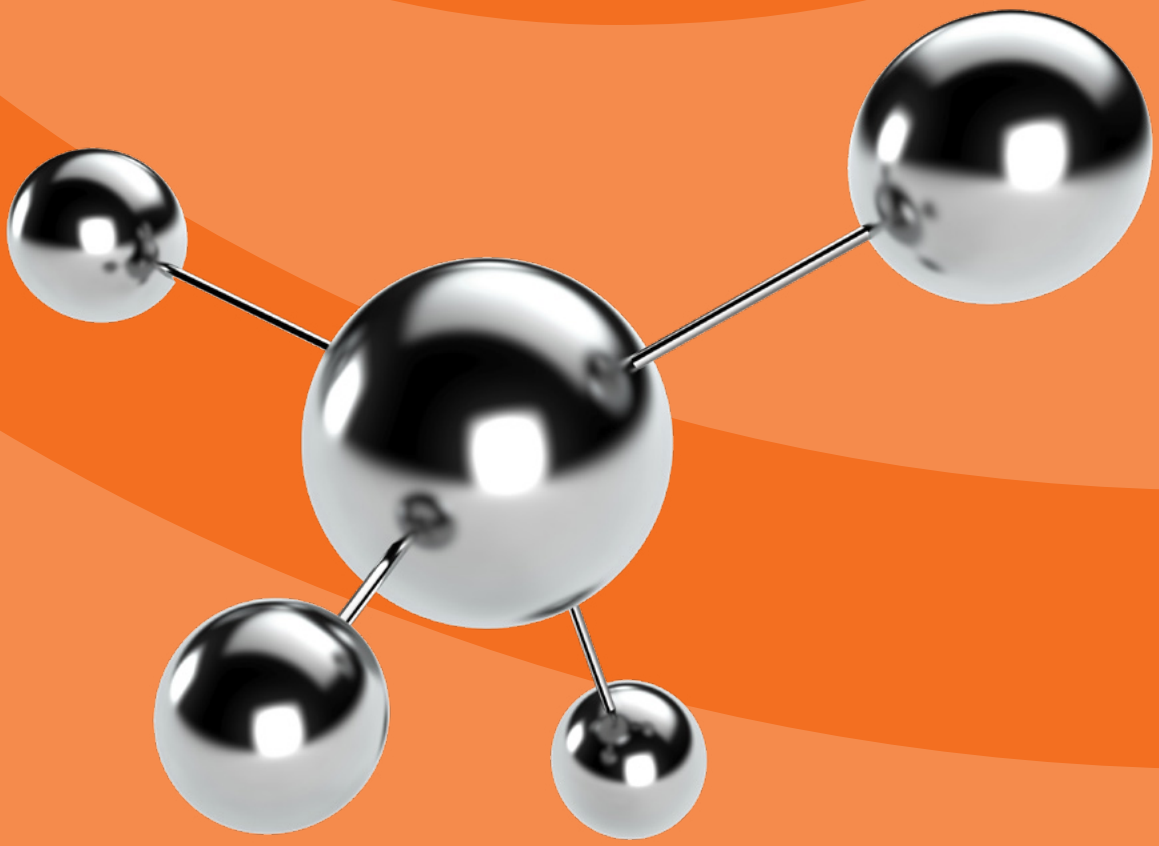

Artificial Intelligence

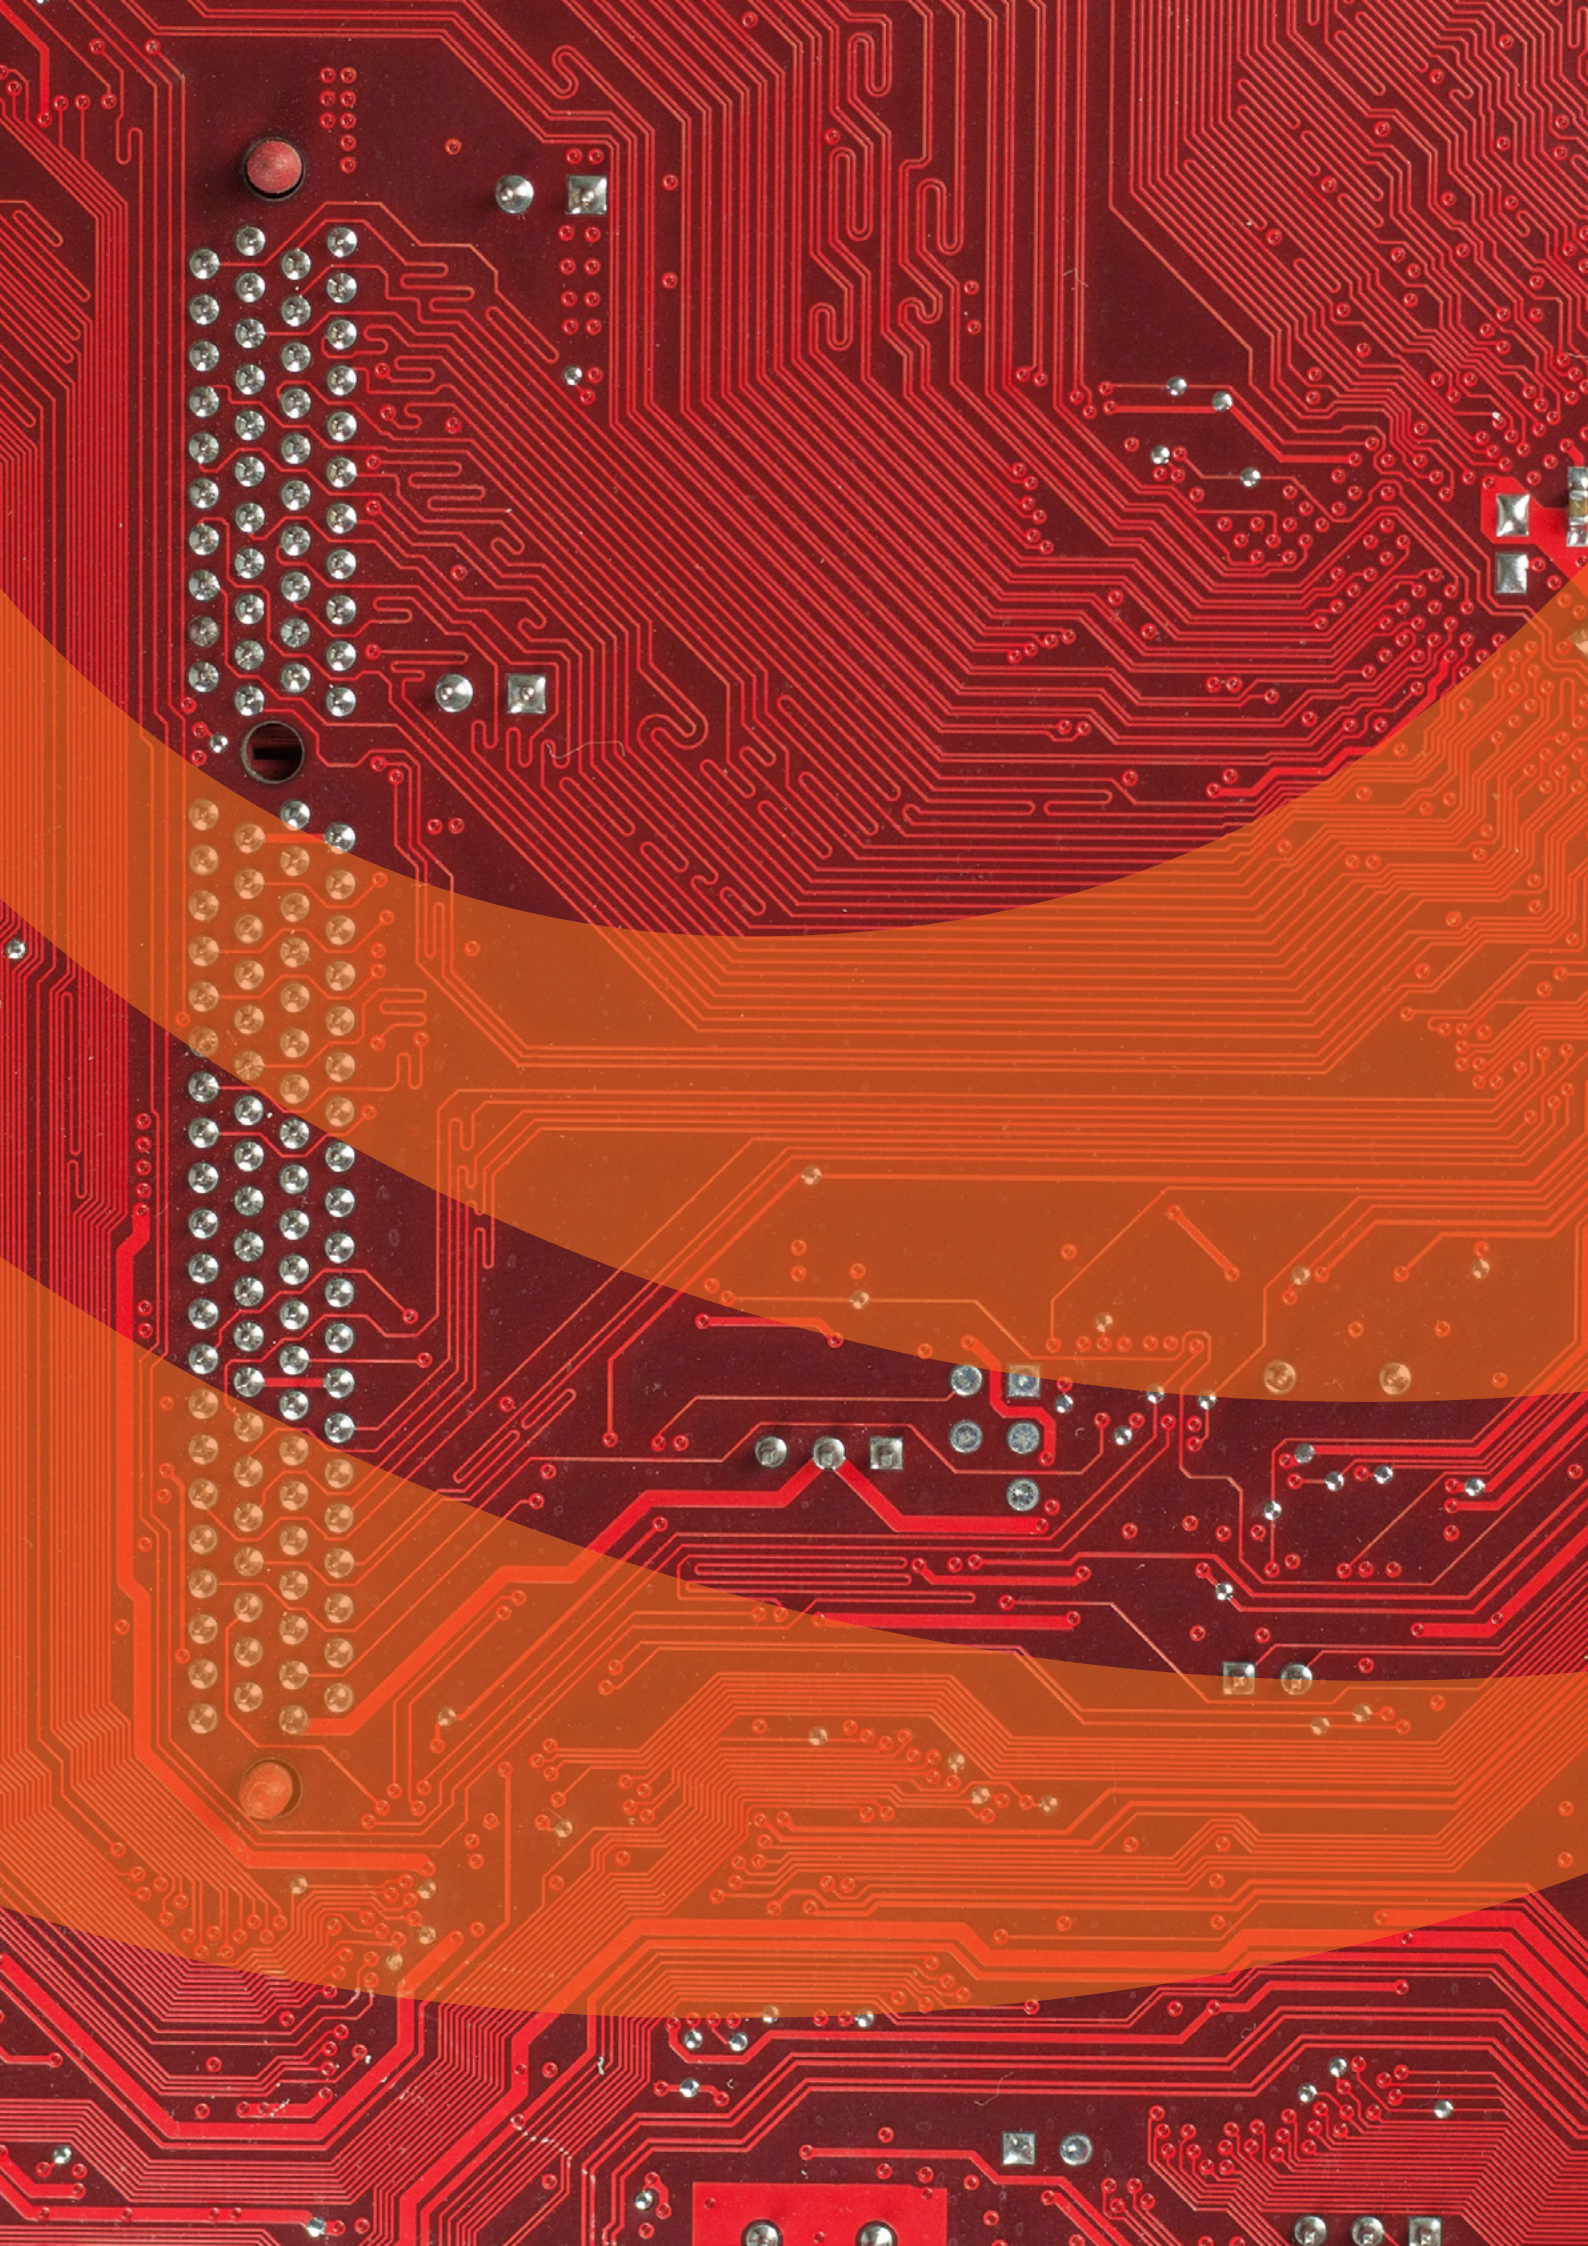

Artificial intelligence is not new: the term itself was coined over 60 years ago. However, the convergence of data volume, processing power and technical capability has convinced many that the AI era has finally arrived. Every industry is looking at ways to use AI to improve efficiency, reduce cost or deliver better products or services, and the life sciences industry is no exception.

As ever when a new technology becomes viable (think of the internet and the Dotcom boom), technological advances race ahead, while we poor lawyers struggle along behind, trying to ensure that laws that never envisaged what has now become – not only possible, but yesterday's news, remain fit for purpose.

At CMS, we have many lawyers in many countries thinking about the increasing impact of AI on the life sciences legal landscape. Rather than keep these thoughts to ourselves, we thought we would share some of them with you. They cover a wide range of issues, so not all may be directly relevant to your work, but we hope that all will help you understand a little better some of the many issues that life sciences companies will have to address over the coming months and years.

We look first at the types of AI that are being developed and how they are being put to use in the life sciences industry. We then turn to some of the difficult legal issues that we and our clients are facing, including a look at the increasingly important Chinese legal landscape, before finishing with a review of "ethical AI" and what AI developers need to take into account to ensure that their products are trusted by consumers.

We hope you find the articles interesting and useful and if you would like to hear more, we would be delighted to continue the conversation with you!

Niall McAlister
Partner, CMS UK
T +44 20 7067 3581
E niall.mcalister@cms-cmno.com

Roland Wiring
Partner, CMS Germany
T +49 40 37630 309
E roland.wiring@cms-hs.com

Defining AI

Artificial Intelligence (AI) refers to the science of developing computer programmes that simulate intelligent (human) behaviour. Since the term was first coined at the Dartmouth Conference in 1956, the field has developed in bursts. The roller-coaster ride is, once again, gathering pace and AI is considered once more to be ready to change our world.

Why has the environment for AI development become so favourable? There are a number of reasons, including the development of much more sophisticated mathematical models. Perhaps the most important is that the processing power of computers has increased to such a degree that it is now possible to interrogate the volume of data necessary to *train* AI programmes to achieve high levels of accuracy. Allied to that, there is now an unprecedented volume of data to interrogate. It is claimed that in 2019 alone, the world will create as much data as in the previous 5000 years. This convergence of data volume, processing power and technical capability has convinced many that the AI era has finally arrived.

Narrow and general AI

So what exactly is AI? At one level, any computer programme that is designed to process information simulates a form of reasoning and could be called AI. Any attempt to distinguish programmes as AI is fraught with difficulty, but the programme should simulate, to some degree, human cognitive thinking and decision-making. Whilst it is commonly accepted that there is no fixed definition of AI, it is also accepted that for a programme to claim AI capability, it must demonstrate at least some of the following behaviours associated with human intelligence: planning, learning, reasoning, problem solving, knowledge representation, perception, motion, and manipulation and, to a lesser extent, social intelligence and creativity.

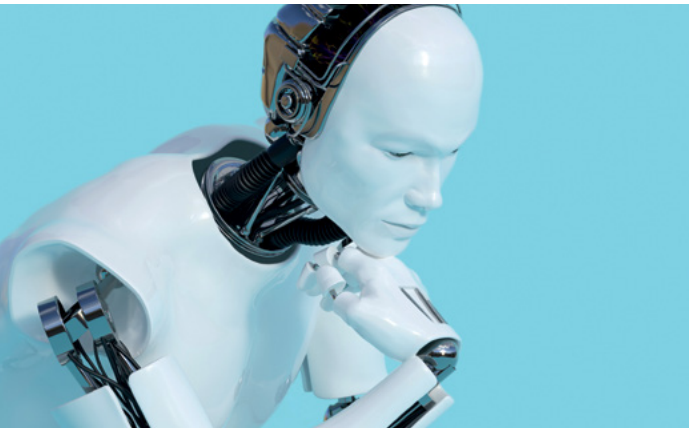

To what extent are existing AI programmes capable of replicating these behaviours? Again, there is no single classification of AI, but one way to categorise programmes is by reference to their increasing level of sophistication. Thus, the most basic forms of AI are reactive machines, then limited memory programmes, followed by theory of mind capability and (the ultimate goal) self-awareness. The first two of these are often referred to as narrow AI and the latter two as general AI. As of today (and most likely for some time to come), only narrow AI has been widely applied in marketable products, as even limited memory AI (such as self-driving cars) is largely in its infancy.

Reactive machines

The large majority of AI products being developed today are (in AI terms, relatively unsophisticated) reactive machines. Insofar as they have any concept of the wider world, it is a very narrow one, and they are not easily reprogrammable to address different tasks. Reactive machines are good at (in human terms) sophisticated games (for example, Deep Blue and chess or AlphaGo and Go), where there is a considerable amount of data to interrogate within a narrow field.

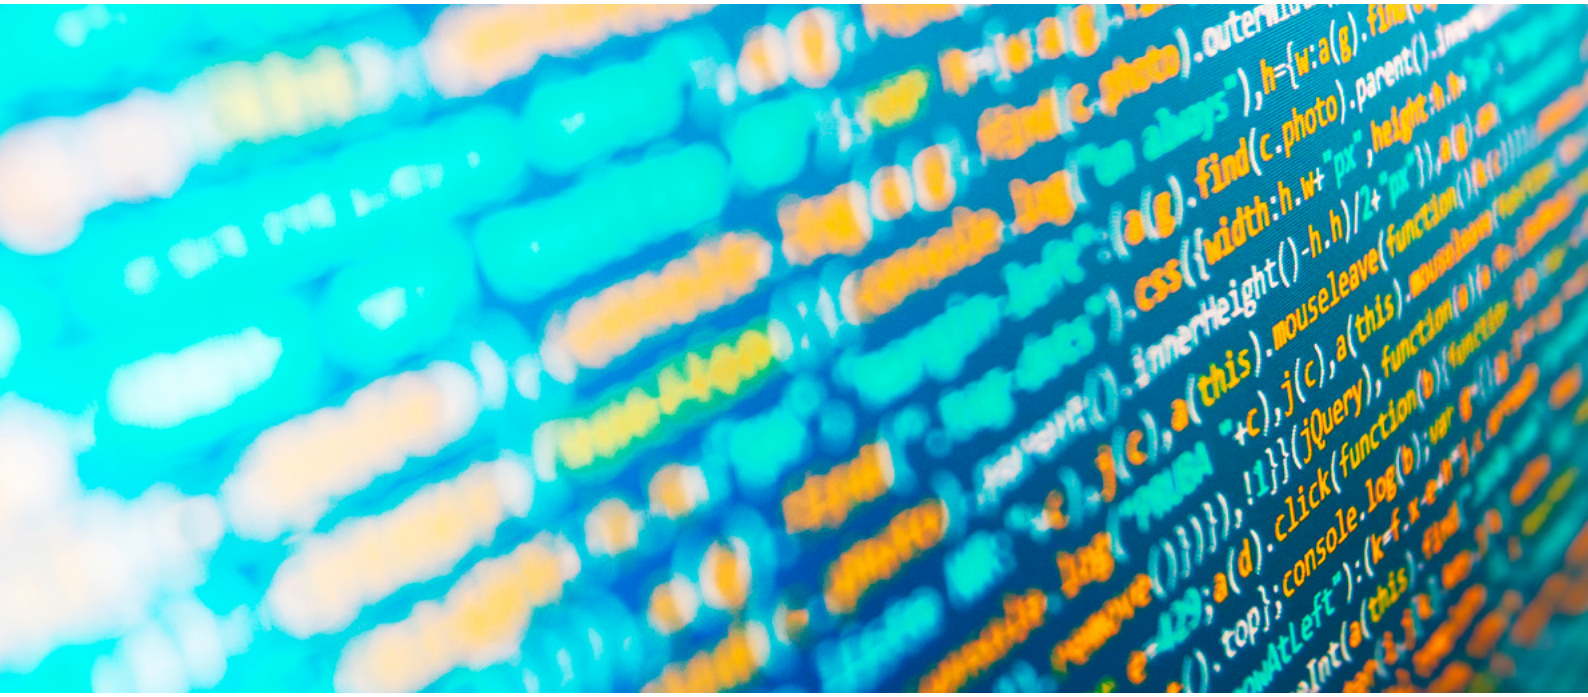

Reactive machines do not create a database of experiences for the programme to exploit creatively. Instead, they apply knowledge to filter out the sub-optimal outcomes. Consequently, each decision Deep Blue makes is a new one based on the chessboard at that time. It does not apply memory or experience to plan moves in advance, as a human would.

Limited memory

Limited memory AI is at an early stage of development. As the name suggests, these programmes are able to make use of historic information. This means, for example, that self-driving cars can retain information about the wider environment they are operating in and apply it to change lane, anticipate another vehicle coming close or a pedestrian walking into the road. However, the information is not retained and used to develop understanding. Beyond its use to interpret the immediate environment, it is lost.

Is AI a single technology?

AI is an umbrella term that covers a number of different technologies, primarily machine learning, deep learning, neural networks, natural language processing and computer vision. There is a considerable degree of connection between them, but the core technology is machine learning.

Machine learning allows computers to learn and improve from experience without being explicitly programmed. The algorithm will analyse data and make a prediction. By doing so across a wide range of data, the predictions the programme makes become more accurate over time.

Deep learning – Whereas machine learning operates on a single layer of information, deep learning is a sub-set of machine learning that allows different datasets and data types to be interrogated in one process. This more closely mimics how our brains work – we will naturally take account of written, visual, aural and other information in identifying a cat or a particular type of car.

Neural networks – Deep learning is made possible by neural networks, which are designed to replicate the neural connections in our brains. Neural networks are not algorithms themselves, but a framework that enables different machine learning networks to interact.

Natural language processing allows computers to interpret, recognize, and produce human language and speech.

Computer vision uses deep learning and pattern identification to interpret the content of an image, enabling computers to identify, process and interpret visual data.

AI – Opportunities and challenges in life sciences

Every industry is exploring ways in which AI can improve quality and reduce costs, and the life sciences are no different. In view of the widespread governmental attempts to rein in the ever-increasing cost of healthcare, the use of AI to reduce the cost of drug discovery and delivery of healthcare services, and to improve the efficacy of product development and speed to market, opportunities are being pursued throughout the industry.

The economics of AI in life sciences

Despite the significant opportunities to use AI to reduce healthcare costs and improve patient outcomes, there are still big obstacles to achieving success. Two in particular mean that the impact of AI on healthcare will be likely to be a slow burn. First, the number of physiological interactions that could or do have a role in any disease or condition, and the many ways the interactions take place, mean that the volume of data that will have to be interrogated dwarfs the volumes necessary in other non-healthcare situations. Consequently, computer processing power and capacity will need to increase even more before complex conditions can be mapped and addressed with confidence. Secondly, the quality of the data affects the quality of the outcome, and much healthcare and medical research data will require considerable curation before it can be used to produce robust answers. That is a time-consuming and expensive process in itself. These issues affect the economics of AI adoption in life sciences. First, the greater cost and high risk of failure will affect return on investment for life sciences companies. It also makes the sector less attractive to the most talented AI developers, who are in very short supply and very high demand. A recent Deloitte report that analysed AI adoption in a range of industries placed life sciences and healthcare in the "High AI investments/low returns" segment (the only industry in that segment).

Impact and potential of AI solutions

Notwithstanding these and other risks, there is a clear drive across the industry to develop AI solutions. We describe below some of the ways in which AI is being developed by (and is considered likely to have a material impact on) life sciences businesses.

Drug candidate selection – screening to identify molecules for drug development is a time-consuming, expensive and inexact process, subject to inherent biases that can cause inaccurate identification. AI programmes are able to interrogate these large, complex datasets more quickly and precisely than before. It can also search scientific literature for relevant studies. It is anticipated that neural networks, which can cross-reference many different information streams and make links that would otherwise be practically impossible, will increasingly provide a much more accurate shortlist of drug candidates more quickly and cheaply than existing processes, which will improve the economics of drug discovery immeasurably.

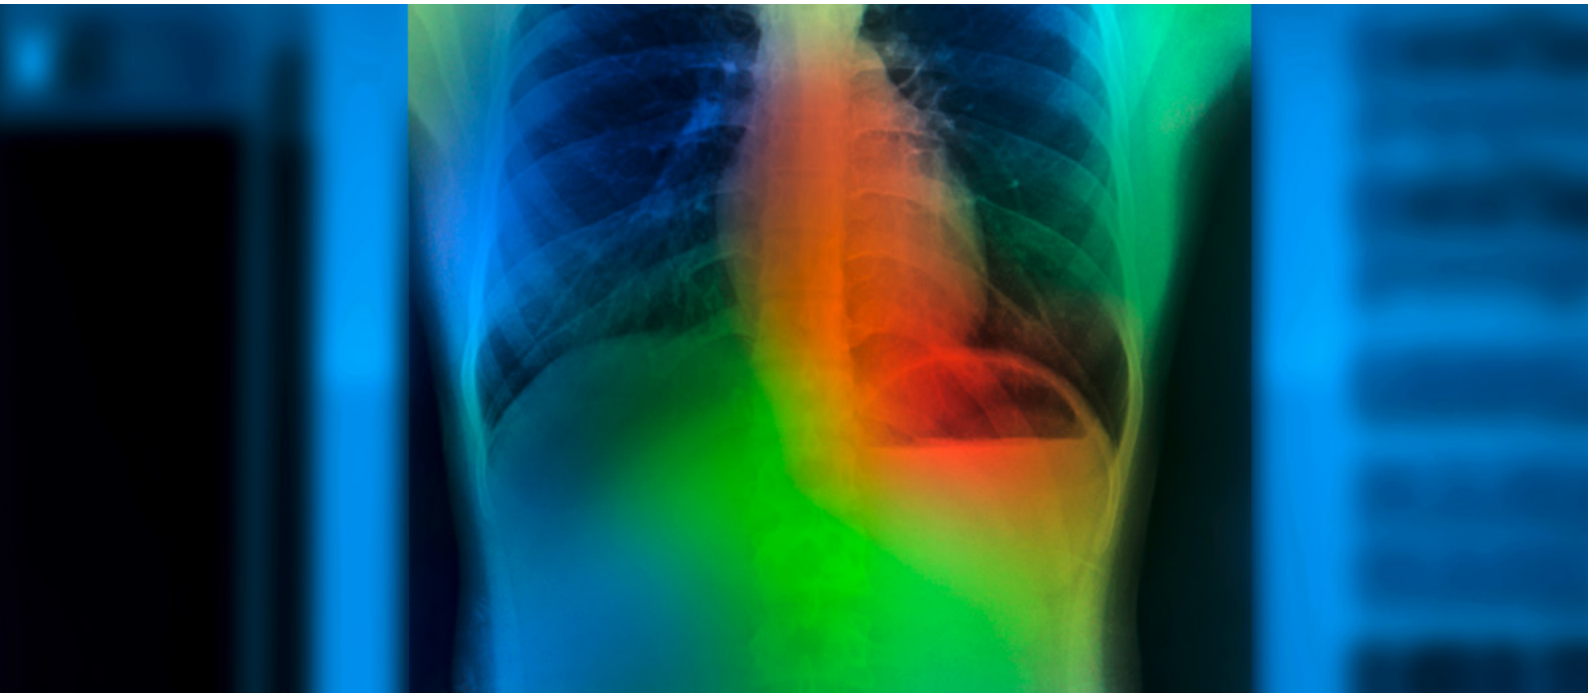

Clinical trial design and data interrogation in drug discovery – for regulatory reasons, late stage clinical trials have been conducted on large, diverse patient populations over many sites. AI is playing an increasingly important role in the design of clinical trials and the interpretation of the data they produce. This is enabling patient enrolment to be much more effective, trials to be conducted on smaller patient populations and for those patients to participate remotely from a wider geographical area, reducing cost while increasing the likelihood of obtaining accurate and relevant data.

Repurposing – AI enhances the possibility of re-interrogating clinical and *in market* data to determine whether existing drugs can be remodelled for other purposes. The incremental cost and the opportunity to repurpose will be attractive economically. There will also be an opportunity to explore data from failed historic clinical trials that may provide new insights into why the candidate failed and which may suggest alternative opportunities that were not obvious before.

Accurate personalised medicine – currently, medicine dosing is relatively generic, in that relatively little information about the individual patient is taken into account when choosing a therapy and setting dose sizes. To a large degree, these decisions are still a matter of trial and error. This will change when AI platforms are able to interrogate the wide range of information about the patient and determine which drug has greatest chance of successfully treating their condition and in what volumes. Also, by continually reviewing data, the platform will, for example, be able to adjust dose size or, if the disease mutates, revise the decision and introduce a more effective alternative.

Patient records – there is little standardisation of patient health records, even where they are digitised. Consequently, it is difficult to extract relevant information or to make connections that may allow meaningful insights into the underlying causes of ill health. AI can overcome some of these limitations. For example, natural language processing tools can ensure that information is captured in a more standardised way, making it more accessible to search tools. Other free text search programmes are able to extract key terms from less structured data. Diagnostic algorithms are helping predict (and therefore track and manage) risk of future illness on the basis of historic health data.

Real world evidence – the increasing availability of real world data as a way to assess performance of drugs in the real world gives rise to a number of benefits. First, it is allowing regulators to approve new drugs sooner – and therefore more cheaply – on the basis that they will be monitored on an ongoing basis for effectiveness and side effects. Secondly, it is enabling healthcare

systems to push manufacturers into payment by results charging models where therapies are paid for on the basis of outcomes. It is expected that this will reduce waste in healthcare systems considerably.

Image recognition – the use of AI to interrogate medical images to identify disease is already under way and is likely to be one of the early success stories in the use of AI in this sector. An example that received a lot of publicity last year was a programme jointly developed by Google DeepMind and Moorfields Eye Hospital in London. The programme was trained on approximately 15,000 images after which it identified eye disease in approximately 1000 images more accurately than a team of retinal specialists. It will be able to complete its analysis more quickly and accurately the more images it reviews. A further benefit is that the algorithm may be adaptable for use in reviewing radiotherapy and mammogram images.

Robotics – robotic surgery has received a considerable degree of attention, particularly the da Vinci robot, which allows surgeons to undertake procedures in otherwise inaccessible places. Not everything robotic surgeons do applies AI but, in one trial, an AI programme outperformed surgeons at suturing wounds in a pig model. Once trained, a robot will be able to perform consistently and accurately no matter how long an operation takes, whereas human performance will inevitably decline with time.

Patient medicine management – patients with chronic conditions currently spend a considerable amount of time, and their healthcare professionals' time, meeting to review their condition and reset doses etc. An "intelligent" applicator could take on much of this process by tracking data about the patient's vital signs and applying tailored doses of medicine, or raising alarms, with only limited input from either the patient or health professional. For example, Medtronic intend to develop autonomous insulin pumps, which monitor patients' blood glucose levels and inject insulin as needed.

Supply chain/logistics – there are many other ways that AI can help transform life sciences businesses. For example, automating processes, forecasting demand and providing insight on collected data are common use cases for AI. Epidemiological AI programmes may help predict where, when and how virulently outbreaks of disease will occur, enabling manufacturers to scale up (or down) production accordingly.

Artificial intelligence and healthcare in China

China is a global leader in artificial intelligence alongside the US and UK and it is the Chinese government that is spearheading this leadership. In 2017, Premier Li Keqiang named AI as an area the government would accelerate progress in, with a particular focus on the life sciences and healthcare sectors.

Focus of public and private sector investment

The government has outlined its ambitions for leadership in AI in various publications:

1. **Made in China 2025** – The Chinese government committed to investing $300 billion in high-tech fields.
2. **Plan for the Development of a New Generation of Artificial Intelligence** – The government forecast that the Chinese AI industry will be worth RMB 400 billion by 2025.
3. **Healthy China 2030** – Through developments in high-tech fields and AI, China predicts their healthcare industry to grow to RMB 8 trillion by 2020 and RMB 16 trillion by 2030.

This investment and ambition is matched by China's private sector. Tencent, the social media giant, is investing significantly in developing "Tencent Doctor" a healthcare app with *predictive* AI capabilities. Fellow tech giant Baidu is following suit with the medical chatbot app "Melody the Medical Assistant."

Law and regulation keeping pace

Despite the clear focus on accelerating development in this industry, local regulations will need to keep up with the rapidly evolving policy landscape. For example, whilst the *Regulations on the Supervision and Administration of Medical Devices* sets strict procedures and standards for the registration and approval of medical devices, the National Medical Products Administration (NMPA) has not yet formulated the registration technical guidelines for AI-related devices, meaning that the relevant approval authorities will not be able to accurately review the technical standards of AI devices filed for registration.

Currently, despite the lack of specific laws and regulations on AI diagnostic services, existing laws permit the use of AI in assisting doctors in the diagnostics process. However, AI software is itself prohibited from being used to provide diagnosis advice independently. Thus, AI chatbots such as Baidu's "Melody the Medical Assistant" would only be able to provide general medical consultancy services, rather than diagnostic services.

A recent flurry of regulations have also been released on remote healthcare to clarify how AI and other digital technologies can be used in medical services. Medical institutions looking to provide "Internet diagnosis" or "virtual doctor" services, via AI or other digital technologies, must possess a practice licence and cannot use these services during a patient's first consultation. The *Administrative Measures for Internet Hospitals (Trial for Implementation)* also requires "Internet Hospitals", being fictitious hospitals operated by and based on real hospitals or medical institutions, to be approved by the relevant authorities before being operational.

Considering the key role data collection and sharing has in developing AI technology, especially in relation to sensitive personal health information, it is telling that there is still no uniform law or national authority providing a regulatory framework on privacy and data protections laws in China. For now, medical health devices that collect information within China are subject to regulation under the *Cybersecurity Law of the People's Republic of China* (Cybersecurity Law) and the various related technical guiding principles. The Cybersecurity Law also introduces the concept of "Critical Information Infrastructure" (CII) operators, a broad term that may include private companies providing cloud computing, big data or other food and drug-related services, whereby storage of all personal information collected during operations within China are to be exclusively kept within the territory.

AI algorithms challenge the life science mindset

Deciding whether to use trade secret protection or patent protection can be a difficult business decision for AI technology owners, in particular as the life sciences sector increasingly uses machine-learning technology.

Neural nets and drug design

Neural nets are now used in medical diagnosis, drug design and biochemical and medical image analysis. In 2015 for example, US company Atomwise introduced AtomNet, a deep learning convolutional neural network for structure-based rational drug design. According to the AtomNet website, the trained neural network model is able to "predict new active molecules for targets with no previously known modulators."

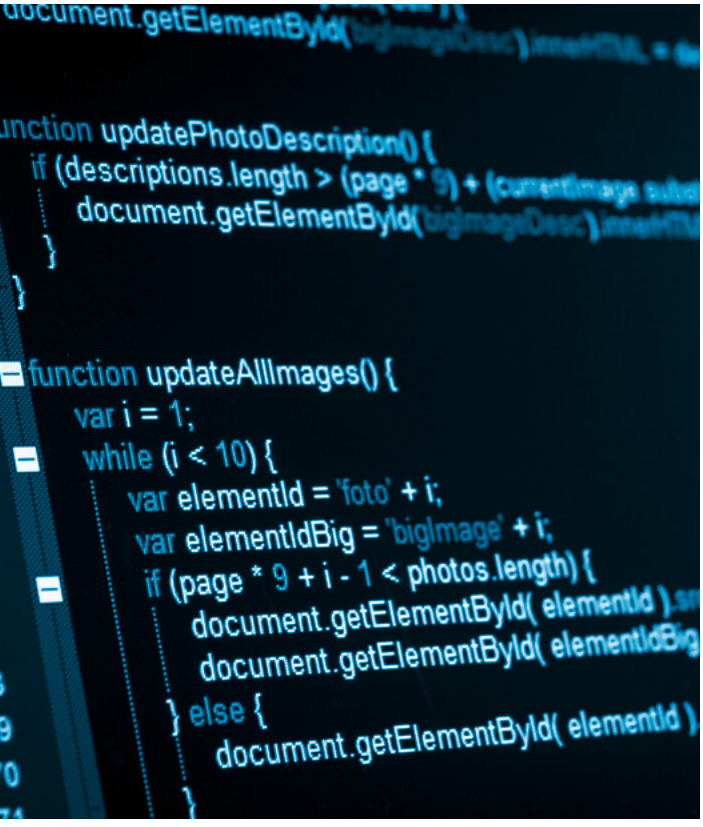

Convolutional neural networks are typically used for image processing and typically take input in the form of an image as a two dimensional array of pixels, each pixel having a red, green and blue channel intensity value. In order to make a convolutional neural network take input that represents a protein-ligand pair, a set of voxels are used, each voxel having channel values for carbon, oxygen, nitrogen and other atom types. The neural network is trained using labelled training data comprising sets of voxels depicting known binding affinities of protein-ligand pairs. It is then able to generalise its knowledge so that when a new potential protein-ligand pair is presented as input, the neural network computes the binding affinity. The neural network can compute predicted binding affinities for huge numbers of possible combinations in order to search for new active molecules for targets with no previously known modulators.

To patent or not to patent

Atomwise decided to proceed with patent protection – the EPO register shows at least three European patent applications. The material is potentially patentable since the purpose of "predicting binding affinity of one molecule to one target protein" is regarded by the EPO as being of a technical nature. In coming to its decision, Atomwise will have had to evaluate the respective benefits and drawbacks of patenting or retaining the programme as a trade secret. The following table sets out some of the matters Atomwise is likely to have taken into account.

Patent Application	Trade Secret
A patent application is typically published around 18 months after filing. Competitors are therefore able to understand the technology by reading the patent publication document.	Trade secret protection relies on keeping the technology confidential.
Patents are registered rights, which give the patent owner a monopoly so that others who independently invent the same technology at a later time are unable to exploit the technology in the territory of the patent.	Trade secrets do not protect against others who independently develop the same technology.
Generally speaking, once a patent application has been filed at the patent office the technology described in the patent document may be disclosed without damaging the potential patent rights.	If the technology is disclosed the trade secret protection is lost.
It is possible to make the technology transparent for ethical reasons.	Transparency cannot be achieved and so there are ethical risks.
Patent applications cost money since it is necessary to pay patent office fees.	Trade secret protection can be very cost effective since there are no official fees to pay. However, there are management and administrative costs to businesses since comprehensive policies and procedures are needed to track and secure trade secrets.
By publishing a patent application about an AI algorithm for finding new drugs, there is a possibility that it is harder to gain patent protection for individual new drugs found using the AI algorithm. This is because the new drugs are arguably obvious since the AI algorithm is known.	If trade secrets are used to protect the AI algorithm then new drugs found using the AI algorithm are more likely to be inventive and so patentable.
If a competitor reverse engineers the patented neural network technology then patents can potentially be used against the competitor. Reverse engineering of "black box" neural networks is possible (see below) where those neural networks are accessed via an API.	If reverse engineering occurs then trade secret protection is lost.
Generally speaking patents last for a maximum of 20 years.	Trade secret protection can last longer than 20 years as long as the technology remains secret.
If patent protection has been obtained there is no need to keep the technology secret.	If the technology needs to be known by several entities, such as software contractors, customers, and a large number of employees, then it may not be practical to keep secret and trade secret protection is not suitable.
Patent protection can act as a deterrent to competitors.	Trade secret protection is invisible to competitors and so does not have the same deterrent effect as patent protection.
Patents are relatively difficult to obtain and the existence of patents therefore provides a metric for investors and others to measure the value of technology.	Trade secrets do not act as a metric in the same way as for patents.

Those developing AI technology in the life sciences sector have a choice between trade secret and patent protection. Sometimes, however, a combination of trade secrets and patents – protecting core technology using patents and peripheral features and extensions to the core technology using trade secrets – may also be a good approach.

AI patentability in Europe

The number of AI-based patent filings has expanded rapidly in recent years, particularly in the United States and Asia. Even in Europe, patent filings grew at an annualised rate of over 50% from 2014 to 2017. To assist inventors, the European Patent Office (EPO) has recently amended its *Guidelines for Examination* (2018 EPO Guidelines), which for the first time include a section relating to AI and machine learning containing advice about how patents related to AI should be assessed.

Method or invention?

As AI and machine learning are *per se* of an abstract mathematical nature, the guidance provided with respect to mathematical methods also generally applies to such computational models and algorithms. The EPO excludes mathematical methods from patentability if a claim relates to a purely abstract mathematical method and does not require any technical means. If a claim relates to a method involving the use of technical means, for instance a computer or a device, the subject matter in its entirety is of a technical nature and is patentable as an invention. The question is then whether the invention satisfies other requirements of patentability, in particular novelty and inventive step.

Technical effect and purpose

The evaluation of the inventive step, widely considered the more problematic requirement, assesses whether the mathematical method contributes to producing a technical effect that serves a technical purpose: For example, an X-ray apparatus providing a genotype estimate based on an analysis of DNA samples or an automated system providing a medical diagnosis by processing physiological measurements.
The 2018 EPO Guidelines, referring specifically

to AI and machine learning, state that when examining the technical character of the subject matter, expressions such as *support vector machine*, *reasoning engine* or *neural network* require close examination as they usually refer to abstract models devoid of technical character.
An applicant seeking patent protection in Europe should therefore establish a causal link to the technical purpose. For example, the use of a neural network in a heart monitoring apparatus for identifying irregular heartbeats makes a technical contribution. The classification of digital images, videos, audio or speech signals based on low-level features, such as edges or pixel attributes for images, is a typical technical application of classification algorithms.

Data and training

The 2018 EPO Guidelines also helpfully specify that generating the training set and training the AI models may contribute to the technical character of the invention if they support achieving the technical purpose. However, classifying text documents solely in respect of their textual content does not serve a technical purpose *per se*, but a linguistic one. Classifying abstract data records or even telecoms network data records does not have a technical purpose *per se* when the resulting classification has no technical use. It is immaterial that the classification algorithm may have valuable mathematical properties, such as robustness.

A friendlier patent landscape?

In summary, AI inventions are subject to the same criteria as any other inventions implemented by computers. The EPO is endeavouring to develop examination practices that are friendlier to such inventions. A good example of this is the recently granted European patent in relation to the use of deep learning for bone segmentation and removal in computer tomography angiography imaging. After the initial decision that patentability requirements were not met, the applicant was able to overturn the EPO's opinion by successfully arguing that the invention provided a method for reliable and precise bone removal in a 3D medical image.

Who owns what when AI does the research?

AI has transformative potential in the life sciences sector, especially when applied to drug discovery research. That research is critical – it solves the unmet needs of patients and provides the funding and income that is the lifeblood of the innovators in the sector. The ability of AI to analyse huge datasets is poised to make a huge difference to drug discovery research. However, the role of AI in drug discovery is not without its problems. Some of those problems are legal problems, including issues around ownership.

Machine inventors

Laws on intellectual property are based on human researchers, human inventors. For example, current patent legislation in the UK, the Patents Act 1977, envisages that patents are granted to inventors. Identifying the inventor was relatively straightforward in 1977. Drug discovery was undertaken by human scientists in traditional labs.

The inventor of a new drug or a patentable invention becomes less clear as the involvement of AI increases. This will be even more so as machine learning improves. If a machine makes the discovery, who is the inventor for patenting purposes? The programmer? The researchers that taught the programme? The company funding the research project?

Companies like Exscientia are already harnessing AI in the drug discovery process. Tech giants such as Google, Apple, Samsung and IBM are moving into the life sciences sector. As AI in drug discovery gathers pace, we need to find practical solutions to this issue.

Data confidentiality

For AI to be effective, it needs useful data to analyse. As the digitalisation of health data increases, so does the amount of useful data available for drug discovery analysis. The catch is that the data is personal to patients. Although the patients do not technically own the data, they do have legal rights in it.

Most types of patient data, particularly relating to diseases, are sensitive. Many patients are uncomfortable with their data being used for anything other than their treatment. The fall-out of the Google DeepMind collaboration with the Royal Free NHS Foundation Trust in London, where the Trust failed to secure the confidentiality of patient data in its contract with DeepMind, highlights the importance of maintaining public trust to ensure continuing access to patient data.

In contrast to the Patents Act 1977, the legislation governing personal data – the EU General Data Protection Regulation and national subordinate laws, such as the UK Data Protection Act 2018 – is recent and was drafted with the digital age in mind. That said, there is not currently a clear process for legally compliant use of patient data in drug discovery.

Keeping up with technology

As safe as a black box in the cloud?

It is often assumed that machine learning models presented as black box services in the cloud are secure; that is, the machine learning model is not revealed to competitors. It is assumed that, without explicit knowledge of the type of algorithms powering the technology, its parameters or the datasets it has been trained on, machine learning models are very difficult to decipher and reproduce.

Machine learning models are often presented as 'black box' services in the cloud, hidden behind user interfaces such as mobile apps, websites, or application programming interfaces (APIs). Given an input, a black box will return an output without exposing the decision-making process hidden within. However, new abilities to reverse engineer these machine learning algorithms are starting to emerge. In a 2016 paper entitled "Stealing Machine Learning Models via Prediction APIs" by researchers from Cornell Tech, the EPFL and the University of North Carolina, research showed that, with nothing more than repeated queries to a machine learning model, it was possible effectively to reproduce a machine learning-trained AI via a *model extraction attack*.

The research team demonstrated it was possible to infer the hyperparameters of a deep neural network hosted on Big ML and Amazon ML on AWS services by observing responses to a sequence of queries sent to an API of a deep neural network service, when the responses included prediction or confidence values associated with the responses. These values allowed the researchers to infer and estimate to a high degree of accuracy the internal parameters of the neural network, effectively *reverse engineering* the neural network and making it potentially reproducible. By using their inference method, any proprietary information or trade secret protection in the original deep neural network was lost.

The law has always struggled to keep up with technology. Now greater challenges face us as the fast moving technology sector meets the heavily regulated life sciences sector. Applying the power of AI to life sciences, and drug discovery in particular, has great potential. The challenge for the law is to keep up with the change and not stifle the benefits.

AI as inventor – a threat to the development of life science AI?

AI already provides an important contribution to drug discovery and development through analysing big data and applying deep learning techniques to analysing large datasets. Patent law in principle allows the protection of the AI technology itself. However, this does not generally apply to AI-derived inventions. Unless and until patent law properly addresses the protection of AI inventions, life science companies might have to implement new business models in order to protect their innovations.

Sustainable innovation

Patent law stipulates human creation as a basis for patentability. Consequently, an invention with a substantial contribution from AI risks not being patentable under current patent law. The usual practice of life science companies to designate a person or a group of persons as inventor for such inventions leaves open the door for invalidity actions against the patents in question.
Some commentators believe that AI will develop sufficiently in future to be itself the holder of rights and that we might regard AI application as the inventor. In the meantime, life science companies face considerable uncertainties about the commercialization of AI driven inventions. How can they secure sustainability of their innovations?

Protection strategies

A possible approach is to seek a change in the law. Several approaches are conceivable. Patent law could establish that the programmers of the AI software are the inventors. The company that funded the development of the AI programme could claim the inventions deriving from it. Alternatively, we might regard the persons feeding the AI with data as the inventors.
Any of these propositions would require a substantial amendment in current patent law. Not only would this be likely take some time to achieve, the consequences of such changes are also not foreseeable.
In the meantime, an alternative approach might be to rely on a broad protection strategy to protect innovations. Life science companies might, for example, focus on the protection of the data generated by AI and commercialise such data instead of merely protecting the invention by way of a patent. Whilst the protection of the AI software is certainly important, so is the protection of the concepts on how AI manages the data processing.

AI and public contracts

Public procurement rules apply not only to conventional purchase contracts, but also to innovation partnerships and other innovative procurement tools. This has implications for companies dealing with tenders related to AI.

A considerable part of the demand for artificial intelligence products and services comes from public institutions and companies that have to make their purchases under public procurement law, such as public hospitals, statutory health insurance funds, universities and research institutions. But public procurement law not only applies to the conventional purchase of items like 3D printing machines for medical devices, mixed-reality spectacles for surgeons or technical equipment for research institutions, instruments with which AI can be used. Other procurement, such as cooperation within the framework of a public-private partnership, joint ventures and think tanks, *e.g.* innovation hubs, can also be subject to public procurement law.

These open up wider possibilities and other tools of innovative procurement, one of which is innovation partnerships. In contrast to traditional types of procedure, innovation partnerships allow contracting authorities to award contracts to several companies at the same time. Moreover, once the research company has provided the required research services, there is no guarantee that the contracting authority will also commission it for the commercial manufacture of the products concerned. The contracting authority could instead terminate the partnership after the research work is completed or refuse acceptance later if the promised cost ceiling or performance level is not met. Companies that offer goods or services in connection with AI need to prepare for this and other innovative procurement tools when dealing with tenders in this area.

Who is in the firing line when AI goes wrong?

Whilst AI may be taking the stage in life sciences, it is also widening the liability pool – the range of people who might be sued – when things go wrong. What happens when an AI system incorrectly diagnoses a patient or incorrectly identifies a disease? If medical professionals use AI alongside their everyday practice and a test associated with an AI system is found to be inaccurate, who is liable? The answer is not clear-cut. All or any of the designers and manufacturers of AI systems, programmers and developers and even users and owners of AI systems risk exposure to claims.

Who owes a duty of care?

In cases of clinical negligence it is, more often than not, straightforward to identify the medical professional, insurer or organisation against whom a claim may be brought. When a patient brings a claim against a medical professional, it is necessary to establish that the medical professional breached a duty of care owed to the patient and that, by breaching their duty, they ultimately caused the claimed outcome. Where an AI system is involved, the medical professional may claim that they relied on the AI product, and that it was reasonable for them to do so. In such circumstances, establishing liability may initially depend on the contract between the medical professional or hospital using the AI system and those that provide the AI system.

Strict liability

It will also bring into play relevant consumer protection and product liability laws. These will often impose strict liability on producers of defective products, in addition to more generally applicable legal remedies such as

breach of contract. Where strict liability applies, the producer is exposed to a considerably greater risk. This can include the manufacturer, the developer and all those mentioned previously.

Who is a producer?

Under the EU Product Liability Directive, a *producer* for these purposes is widely defined. It can include:

- the manufacturer of a finished product;
- the producer of any raw material;
- the manufacturer of a component part of a finished product;
- any person who by putting their name, trade mark or other distinguishing feature on the product presents himself as its producer;
- any person who imports a product into the community; or
- a supplier of the product where the producer cannot be identified.

Accordingly, the designers and manufacturers of AI systems, programmers and developers and even users and owners of AI systems could all fall within the scope of this definition. It is also worth noting that whilst all of these parties are potentially in the liability pool, the inclusion of one does not necessarily relieve another. In a recent English case concerning an alleged defective medical device, it was noted that the fact that there was a professional healthcare intermediary present did not provide a complete or automatic defence for a producer of a medicinal product. The European Commission expects to issue further guidance on the Product Liability Directive later this year, in which it will consider in greater depth liability in relation to artificial intelligence.

Legal implications of AI in personalised health care

The most significant advantage of AI is its capability of making sense of large amounts of data in seconds, recognising patterns and structures, from which they then infer certain rules. Based on these rules, AI is supposed to draw assumptions and predictions. It is assumed that these features of modern AI will bring vast progress to the field of personalised medicine. AI will certainly cause massive changes in the health care sector within the next 10 to 20 years. Lawmakers should watch carefully and provide for legislation that fosters the appropriate application of AI in the health care sector while at the same time implementing high ethical standards.

The promise of progress

Employment of AI is particularly promising in functional precision medicine, where millions of human genomes can be sequenced and screened for mutations to identify and develop effective drugs. Before the advent of AI, it was not possible to utilise such huge amounts of data properly.
AI has already quite successfully been tested and employed in major disease areas such as cancer, neurology and cardiology. In stroke, for instance, AI applications utilising the data generated by wearable devices have achieved a >90% accuracy of recognising the movements of patients while experiencing a stroke. Furthermore, AI has been used to predict the outcome and analyse the performance of stroke treatments.

Questions of liability and professional conduct

However, with new methods and applications, there are novel legal issues too. The most predominant: What if the AI-powered assessment is wrong or the prediction is inaccurate? The answer to this question is, of course, complex. First, we must ask if the AI will subsequently make an automated decision based on the result it generated. For instance, could certain medications be automatically administered due to a patient's behaviour, which the AI interprets as the result of a medical condition? Or will a doctor look at the result of the AI and weigh it against his or her professional opinion? While the former scenario is advantageous in terms of response time, it certainly raises complex legal issues, most importantly, who will be responsible for damages that occur based on inaccurate results or predictions of an AI. Machine Learning and Deep Learning algorithms naturally and inevitably produce inaccurate output too (so-called false

positives/negatives). In particular, should we apply strict liability as in product liability law? This would certainly mean a great disincentive for implementing automated decision making in health care. Or should we remain with the traditional tort and contract law approach, where we look for fault or negligence? If so, we will have to ask ourselves how to handle cases where there is no human at fault. And to which standard of accuracy should the AI be held to make decisions without human intervention – the same or a higher accuracy than the average accuracy of a doctor? Another question is whether we need to revise the laws of professional conduct in the healthcare sector. To date, in most countries doctors must practice their profession in a personal and direct manner, an obligation that clearly contradicts automated decision making in health care. Going forward, we will have to define how much human intervention, if any, is still required.

Standard treatment or unlicensed drugs?

Also, off-label use of drugs or use of drugs for which a marketing authorization has not yet been granted may be facilitated and promoted by AI. For example, an AI may identify an unlicensed drug as the most suitable medication to treat a certain patient or disease. Under current laws use of unlicensed drugs is usually only permitted in exceptional cases as a last resort when all other treatment options have failed and under the sole responsibility of the treating physician. Does that mean that the results of the AI have to be ignored and the patient needs to get standard treatment and only if that fails the recommendations of the AI may be followed? As the law stands, it would appear that standard treatment needs to be administered first, even if the AI clearly advises against it. This may also call for a change in current laws.

What happens when smart machines make mistakes?

The implementation of AI systems in healthcare and life sciences is creating a growing benefit for patients. At the same time, it raises new questions regarding product liability in case of damages caused by those systems. The use of AI systems in the development of drugs, diagnosis of patients and the design of individually tailored treatments will make it necessary to rethink product liability principles as well as the split of liability between manufacturers, physicians and patients.

A challenge to the product liability regime

AI systems have reached a level of complexity that confronts the current product liability regime with certain challenges:

- The manufacturer or physician working with the system might not be able to understand all the reasons for the outcome of a certain diagnosis or action. What are the consequences for the determination of liability?
- As a result of the cooperation between physicians and (multiple) AI systems the spheres of responsibilities become increasingly blurred. AI systems work together, exchange data and influence each other's output. How can causality of actions be determined under those circumstances?
- Constant progress of AI systems by deep learning makes it difficult to determine the state of the art at the time of the product's release and therefore to define a product defect in the first place. Do we need a new product defect definition?
- From the damaged party's point of view this also creates a problem to show the necessary evidence. How can new technologies help with that?

Potential legal solutions

The implementation of new unknown medical solutions does naturally go hand in hand with certain risks. The role of a product liability regime along with regulatory laws is to balance the benefits of new treatments and products with the potential risks for patients and on the other hand split the remaining risks among the persons involved. The current legal discussion follows that basic line, although there is a debate if those challenges can be solved with the existing product liability regime or if new specific regulations are needed.

There are two main approaches for the determination of liability with regard to AI systems. One could concentrate the risk with one person and determine it by reshaping prerequisites and definitions. Alternatively, one could also try to create a system of shared risk among all participants. There are interesting arguments and solutions for both approaches which range from implementing a system of mandatory insurances to more exotic ideas like directing claims against the AI system itself by creating a new type of legal person to be held (financially) responsible.

Another solution might be to go one step back and define additional regulatory rules with regard to testing requirements and risk analysis prior to the launch as well as strict monitoring duties for AI systems on the market. The concomitant implementation of blockchain based "Black Boxes" may make it easier to determine causality and at the same time help claimants to present the necessary evidence.

The European Union has reacted to those developments by initiating an expert commission which is currently working on concrete suggestions. A first paper has been published in December 2018 and will be followed by a report in 2019. The next months will therefore most likely bring some interesting movement in the current discussion.

Medical software based on AI – Stretching the regulatory boundaries

The disruptive potential of AI presents new challenges to the systems that regulate medical products and services. The current regulatory framework often does not adequately reflect the rapid developments occurring in the sector. In practice, industry players need to apply the existing framework to these new developments. Stakeholders need to work together to ensure that the regulatory systems evolve in time to embrace the future benefits of AI.

Medical device classification

A central issue in the development of software solutions based on AI is the regulatory classification of such software. Does the software qualify as a medical device? This is of practical importance because medical devices may only be placed on the market if they bear a CE mark and have been checked in a conformity assessment procedure. If a product qualifying as a medical device is distributed without a CE mark, competitors may request the distribution to be discontinued. Placing such a product on the market may even constitute an administrative offense or have criminal consequences, like in Germany, for instance. According to the European Council Directive 93/42/EEC concerning medical devices and according to the new European Regulation (EU) 2017/745 concerning medical devices (MDR), which will become effective in May 2020, software solutions may qualify as a medical devices. The qualification depends on the intended purpose of the software. Broadly speaking, if the software detects or helps to treat illnesses, for example by supporting the diagnosis via an image recognition or by calculating the dosage of medication, it will likely qualify as a medical device. If the software only provides knowledge or only saves data, the product will likely not have an intended medical purpose.
The MDR will bring about additional challenges for industry players developing medical software involving AI. It contains a new specific classification rule only for stand-alone software. According to this rule, most AI solutions will likely be upgraded from class I to at least class IIa. This means the conformity assessment procedure may no longer be carried out in-house by the company itself, but must involve a notified body, i.e. an external auditor. This will increase the administrative burden considerably. Some argue it may even discourage the development of such solutions within Europe. In any event, a company developing software solutions should consider this aspect very early on in the development process.

Professional medical codes

Medical devices involving AI that interact with healthcare professionals may blur the thin line between providing a value-added service to

support clinicians and itself providing a medical treatment service. This can conflict with the applicable rules of the medical profession. An example is surgery support software, surgery robotics. If software is designed to help the doctor during surgery, for instance by suggesting how concrete steps of a surgery are carried out, the question arises to what extent it is still the doctor carrying out his medical profession. Is it in fact the software and the robot carrying out, or at least considerably influencing, the course of the surgery.
This is relevant because in most jurisdictions the provision of medical services is reserved to healthcare professionals, namely doctors. To avoid crossing this line and being exposed to legal and compliance risks, an analysis of which functionalities could run foul this principle and how infringements can be avoided, should happen early in the software development phase

Reimbursement of AI applications

Could a healthcare system reimburse software that analyses magnetic resonance imaging (MRI) and provides a reliable cancer diagnosis without the involvement of a doctor? Would a public sickness fund pay a hospital or a clinic for using such software, even if no doctor were involved in the analysis and diagnosis? These are complex questions to which the answers may differ from country to country. The legal framework is in constant flux. From a business perspective, this is of crucial importance because it obviously has a considerable impact on the financial prospects of an AI based development. Putting this legal and regulatory aspect on the project agenda early on is essential for the success of a business idea based on AI.

Regulatory hurdles and openings

The examples above show that the regulatory framework puts important constraints on the development and use of AI applications in practice. Regulatory hurdles should be identified as soon as possible in the development process of AI projects to avoid hick-ups at a later stage. At the same time, uncertainties in the regulatory framework also offer important opportunities. Using the regulatory leeway can help pave the way for unknown AI applications benefitting patients, industry players and healthcare providers alike.

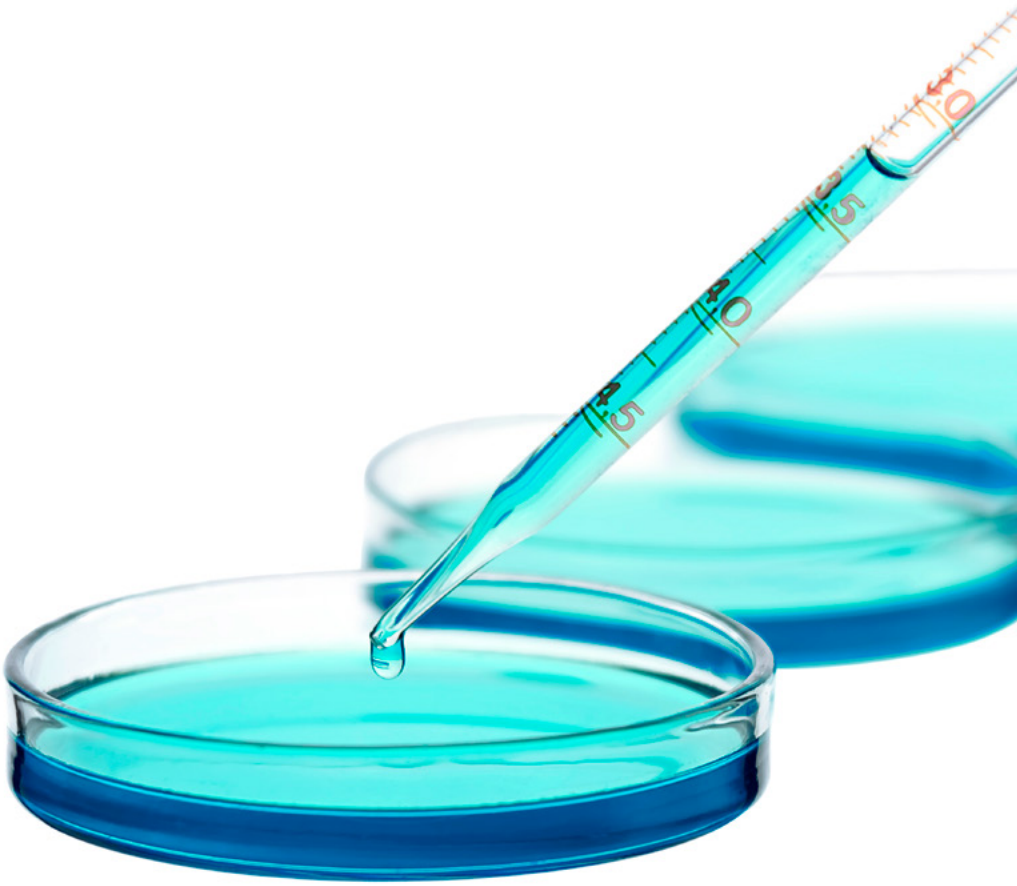

Ethics and AI

The rapid uptake of products that incorporate artificial intelligence across all business sectors has not been met with universal approval. For many, the social impact of the changes they are expected to unleash are an afterthought, at best. However, these social impacts will be real and, if not properly addressed, may have a significant impact on the speed of adoption of new AI-enabled technologies. As a result of these concerns, a new term is becoming more widely used: ethical AI.

Ethical AI is, conceptually, AI that is developed to be socially useful and also socially responsible. It is in the interests of those who are developing AI products or planning to use them in their businesses that they first consider the consequences of their use of AI.

A higher evidential standard

Businesses will need to address a number of ethical issues. Perhaps the most fundamental issue is how far information provided by AI enabled processes – the outputs of AI – should be taken into account in making healthcare decisions and diagnoses. Currently, healthcare practitioners take account of information from a variety of sources (including from the patient) and then, using their own knowledge and experience, make a decision. Much of that information is not of a technical nature. To what extent should supposedly objective AI-derived data take priority over information provided by the patient or, if it points to a different treatment than the practitioner would have recommended, to the practitioner's own judgement? A concern is that practitioners will defer to the AI-derived data because the risk of overruling it would expose them to censure, or that patients will assume it is more valid than their own interpretation (or the practitioner's expertise) and feel that "it must be true" and therefore must be followed. Accordingly, it will be of great importance that AI-derived information is not assumed unthinkingly to be the "gold standard", but that practitioners continue to use it only as part of their decision-making process. This will mean holding AI derived information to a higher evidential standard than we would require of a practitioner, at least until we have a high degree of trust in the AI programme. This makes sense if we accept that the AI programme is intended to reach a factually objective conclusion, whereas we expect a practitioner to exercise judgement in reaching a decision.

Informed consent

Another important issue is the protection of individuals' anonymity when their healthcare data is used, for example, to train algorithms. The question of ownership of data (is it the individual's or the entity that captured it?) and who can have access to it, is unresolved in many jurisdictions. It is becoming increasingly necessary for individuals to have to grant *informed consent* to the use of their sensitive data. The concept of data privacy is not unique to its use in AI, but because of the large datasets required to train AI algorithms, it presents particular problems in ensuring that the algorithms are robust.
It may be practically impossible to obtain consent for historic datasets, so should

consent be dealt with differently where the data is historic? Further, is there a wider public benefit that means there is a presumption of consent for health data unless withdrawn by the individual? This may be easier to impose in countries that offer widespread public healthcare services as a *quid pro quo*, but to what extent should those using the datasets be required to ensure anonymity or be allowed to profit from the outputs of the research without offering a return to the healthcare system? A number of European governments are currently grappling with this issue.

Pitfalls of data bias

The question of data bias is one that looms large in any consideration of the validity of AI outputs. This takes different forms. First, the algorithm itself has to be designed in a way that means it accurately interprets the data it reviews. This is a lot harder to do than it sounds – every software developer has their own inherent assumptions and biases and these can cause distortion in the programmes they write. Secondly, the quality of the data, the volume of the data and the nature of the data can all create bias. For example, the overwhelming majority of clinical trials take place in Europe and the USA and – even within their populations – middle class Caucasian males are significantly over-represented in those trials. Consequently, racial genetic variances, socio-economic influences, or gender-based physiological differences are already not properly taken into account in such trials. Datasets compiled from such trials will carry forward these biases unless they are properly curated and, if they are not corrected, the algorithms they train risk becoming increasingly inaccurate as they develop from an unrepresentative base.

Transparency in complexity

A related issue is the question of transparency. Many owners are likely to treat AI algorithms as trade secrets, rather than trying to patent (and therefore publicise) them. As algorithms are expected to evolve as they are trained, the ability to discern the logic they apply in their interpretations quickly becomes impenetrable. Without transparency, there is the risk that regulators may, for example, discount trial outcomes, or that patients or practitioners will not fully trust the programme's decisions or diagnoses.
One solution may be to enable the programme to explain its decision but, whilst that may assist in relatively straightforward scenarios, the complexity of a programme that assesses a wide range of diverse health data before coming to a diagnostic decision may not be able to do so in a way that even expert practitioners would understand.

AI ethics boards

Questions of trust, safety and transparency are already high on the list of issues that life sciences companies know they must address, and the industry engages with patient populations more extensively than companies in other industries do with their customers. Accordingly, there are already structures in place that can facilitate the ethical development of AI in this sector. Indeed, some companies have already established AI Ethics boards. These boards will need to be able to take account of a range of inputs – as well as therapeutic expertise and a technical understanding of the AI products, they will need legal and regulatory specialists, and, ideally, members with an understanding of ethics who will be able to ensure that the framework remains relevant as the field develops. One issue in the structuring of such boards is whether they should be accountable only internally (to the board of directors of the company), or whether they are also seen to be accountable to the public, such as by publishing an annual report of their activities. In addition, they need to have sufficient authority to make recommendations and to critique their company's approach to AI adoption and use, to ensure that they do actually influence their companies' policies. They are likely to be an important way to build trust with both the public and healthcare professionals.

Contributors

Nick Beckett
Managing Partner, CMS China
T +86 10 8527 0287
E nick.beckett@cms-cmno.com

Roland Wiring
Partner, CMS Germany
T +49 40 37630 309
E roland.wiring@cms-hs.com

David Dennis
Senior Associate, CMS UK
T +44 141 304 6088
E david.dennis@cms-cmno.com

Dr. Sven Brockhoff
Counsel, CMS Germany
T +49 711 9764 164
E sven.brockhoff@cms-hs.com

Jessica King
Associate, CMS UK
T +44 114 279 4251
E jessica.king@cms-cmno.com

Dr Rachel Free
Partner, CMS UK
T +44 20 7067 3286
E rachel.free@cms-cmno.com

Alexander Leister
Counsel, CMS Germany
T +49 711 9764 208
E alexander.leister@cms-hs.com

Niall McAlister
Partner, CMS UK
T +44 20 7067 3581
E neil.mcalister@cms-cmno.com

Roxie Meng
Associate, CMS China
T +86 10 8527 0259
E roxie.meng@cms-cmno.com

Valentina Parvu
Senior Associate, CMS Romania
T +40 21 407 3825
E valentina.parvu@cms-cmno.com

Klaus Pateter
Attorney-at-Law, CMS Austria
T +43 1 40443 6850
E klaus.pateter@cms-rrh.com

Gabriela Staber
Partner, CMS Austria
T +43 1 40443 4850
E gabriela.staber@cms-rrh.com

CMS Law-Now™

Your free online legal information service.

A subscription service for legal articles on a variety of topics delivered by email.
cms-lawnow.com

CMS locations:
Aberdeen, Abu Dhabi, Algiers, Amsterdam, Antwerp, Barcelona, Beijing, Beirut, Belgrade, Berlin, Bogotá, Bratislava, Bristol, Brussels, Bucharest, Budapest, Casablanca, Cologne, Dubai, Duesseldorf, Edinburgh, Frankfurt, Funchal, Geneva, Glasgow, Hamburg, Hong Kong, Istanbul, Johannesburg, Kyiv, Leipzig, Lima, Lisbon, Ljubljana, London, Luanda, Luxembourg, Lyon, Madrid, Manchester, Mexico City, Milan, Mombasa, Monaco, Moscow, Munich, Muscat, Nairobi, Paris, Podgorica, Poznan, Prague, Reading, Rio de Janeiro, Rome, Santiago de Chile, Sarajevo, Seville, Shanghai, Sheffield, Singapore, Skopje, Sofia, Strasbourg, Stuttgart, Tirana, Utrecht, Vienna, Warsaw, Zagreb and Zurich.

---

cms.law