Charles
Russell
Speechlys

Artificial Intelligence *Business Guide*

Get started >

Charlesrussellspeechlys.com

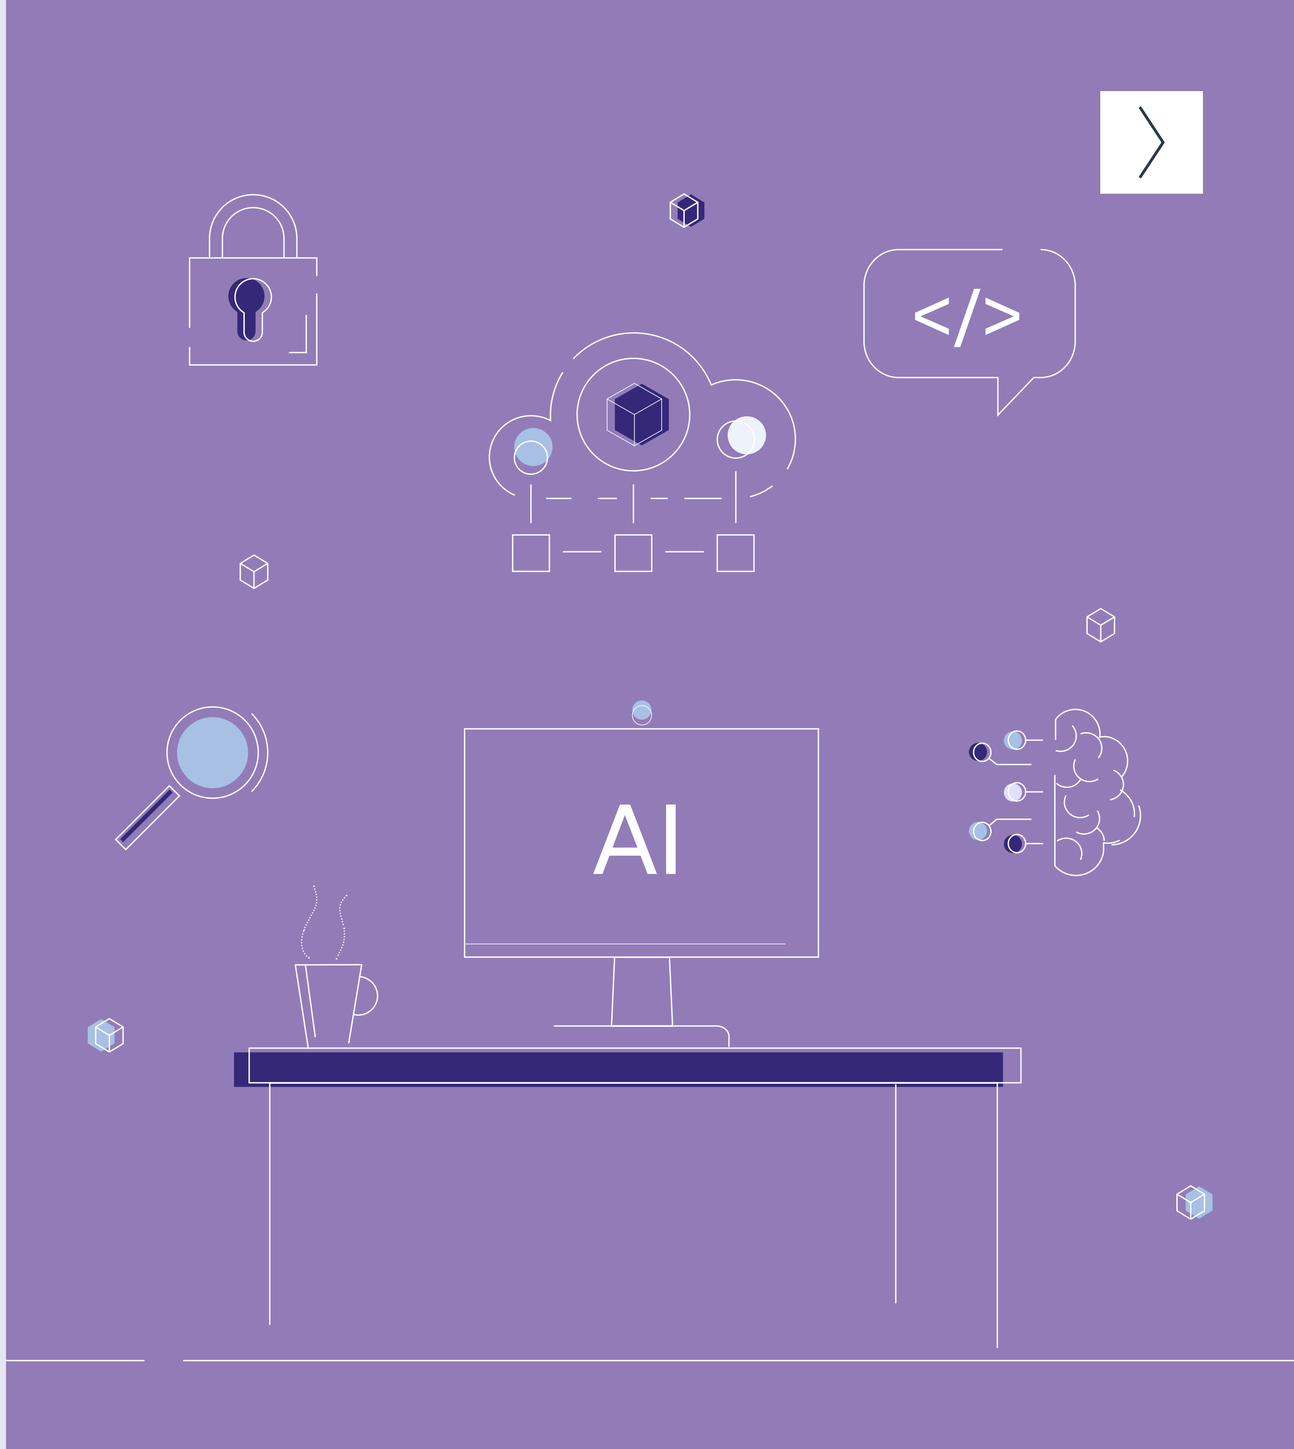

Introduction

The AI revolution is upon us, and it is transforming every sector, including law.

Early forms of AI, such as machine learning, have been with us for a while but the advent of generative AI has made us all recognise the power of the technology and the potential risks of unregulated or poorly controlled use. The legal issues associated with this immensely powerful technology are complex, because specific new laws and regulatory models are being implemented, supplementing existing laws which already apply to AI.

Our key aim in providing this guide is to aid in-house lawyers and senior executives in improving their understanding of AI and some of the key legal issues associated with it. This has involved contributions from practice area experts from across the firm to provide a perspective on the legal implications of AI and how the rapid development of different types and branches of AI creates additional and more complex issues that require consideration.

Our sincere hope is that this guide will empower its readers to spark conversations and identify opportunities and risks in adopting or creating AI based products. We hope to encourage forward-thinking strategies and a business culture based on effective governance and responsible and innovative AI use within organisations.

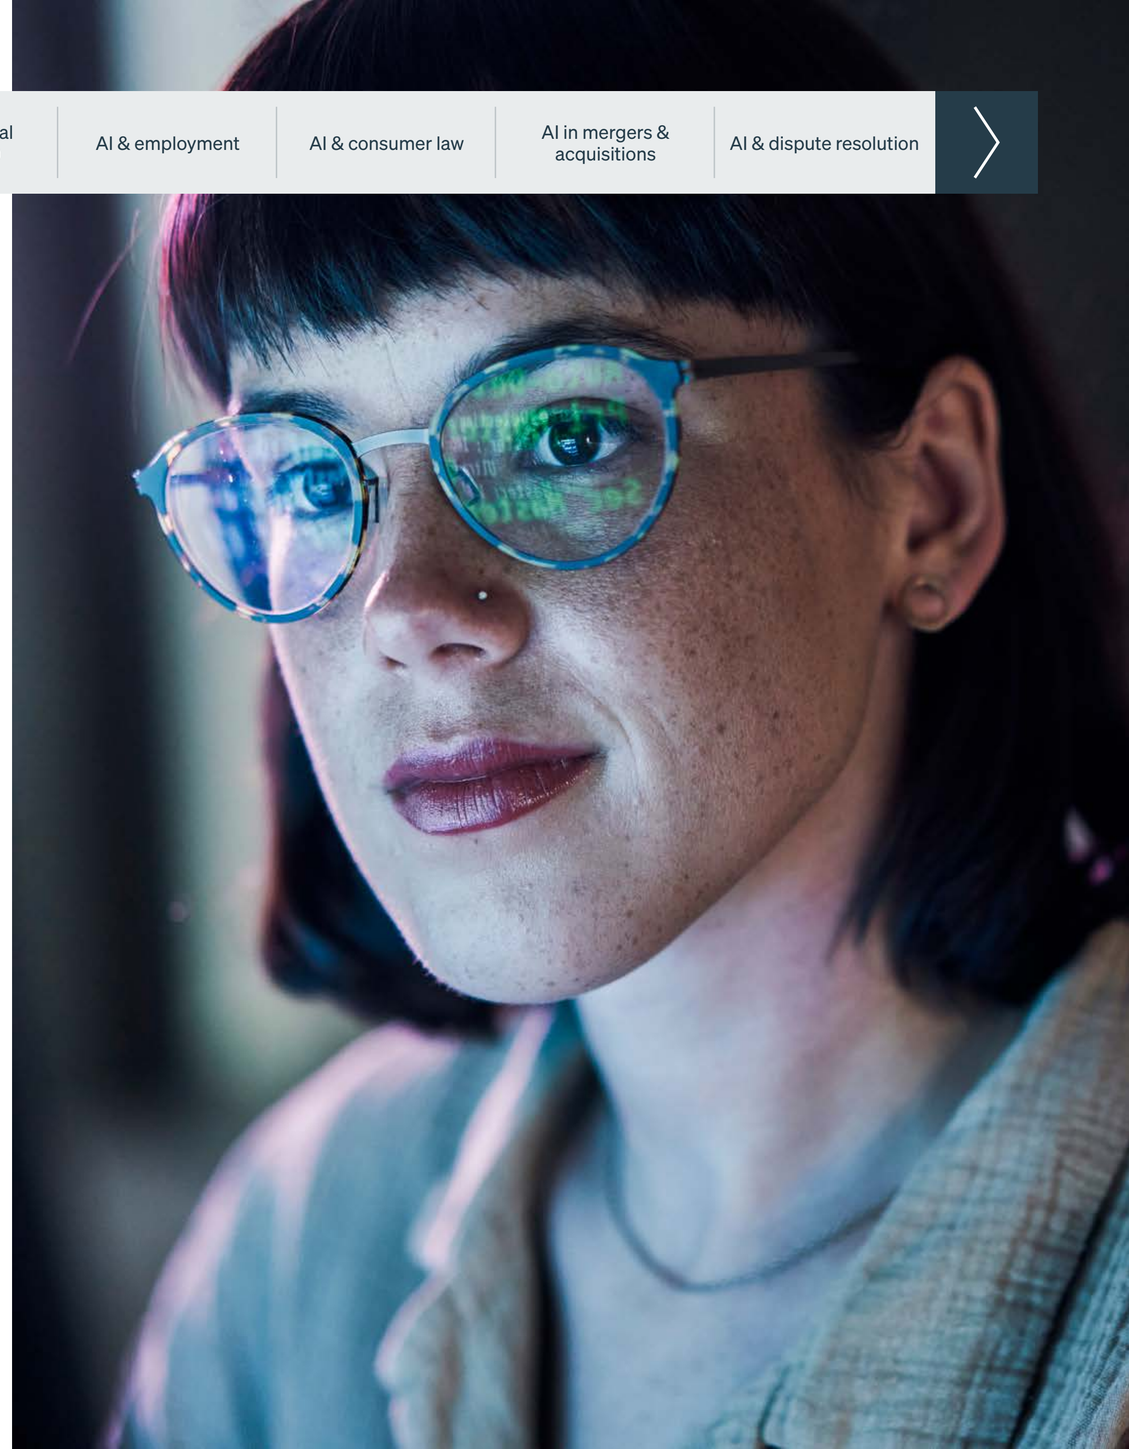

AI regulation and ethics

Governments and international organisations around the world, including the OECD, are looking to be leaders in the rapidly developing field of AI. This involves a difficult balance between having a pro-innovation regulatory regime whilst having sufficient oversight, principles and technical standards (overlaid with robust legislation) to address potential harms and public distrust. In the UK, the government released its AI White Paper in March 2023, updated in June 2023, which outlined its position.

How is the UK government proposing to regulate the world of AI?

The government is proposing a pro-innovation framework underpinned by 5 principles addressing

- safety, security and robustness;
- transparency and explainability;
- fairness;
- accountability and governance; and
- contestability and redress.

Initially these principles will be addressed on a non-statutory basis and implemented by existing industry regulators such as the Information Commissioner's Office and the Intellectual Property Office. The government's view is that rigid legislative requirements at this stage could hold back innovation. At the same time, it is proposing a central oversight function that will help co-ordination across sectors and which may propose additional statutory duties and legislation in the future if this is seen as necessary.

What is happening in the EU?

The EU is leading the way in developing AI specific laws that will apply across all sectors. This contrasts with the currently proposed non-legislative approach being taken in the UK. The AI Act (which sets out AI prohibited practices) and the AI Liability Directive (intended to facilitate claims in relation to AI) are yet to be implemented in the EU but anyone doing business in the EU, or using AI affecting EU citizens, will need to keep an eye on their progress and comply with them when they become law.

What do I need to be doing now and what to look out for in the future?

At present, although the specific AI laws or regulatory regimes are not yet in force, there is significant existing legislation and regulatory guidance that already applies to the world of AI. Data protection, non-discrimination and consumer legislation including product liability, (referred to in more detail

AI regulation and ethics (cont)

in the sections below) are all highly relevant in the world of AI although not necessarily easy to apply in practice.

An immediate priority for businesses is to consider their own internal policies regarding use and development of AI to secure compliance with both current legal requirements and regulatory guidance. Businesses should also, as far as possible, anticipate how the proposed specific AI laws give workable help to understand how businesses can adopt and create AI in a safe, non-discriminatory and ethical manner. Going forward, individual industry regulators will be consulting on and issuing more specific AI guidance to help address this novel and rapidly developing area. Despite the government's initial stance, if guidance alone is not enough to address current or newly identified harms, more robust legislation may follow here in the UK.

No doubt the government will be keeping an eye on how this develops at an EU and indeed worldwide level.

Laying the groundwork for the use of AI

There are three fundamental considerations when evaluating potential use of AI: (i) data (including data inputs, data within the AI model and the outputs); (ii) the AI model and (iii) the governance necessary to provide sufficient guard rails for the development and management of AI. All these considerations must be evaluated to assess the risk associated with a specific AI solution and how to manage it safely.

The specific legal principles in this guide will also be relevant to assessing the risks of developing and acquiring AI within this framework.

Whatever data is applied within an AI model, businesses developing AI must follow recommended practice on conducting appropriate risk assessments, including an AI risk impact assessment, to identify and mitigate risk.

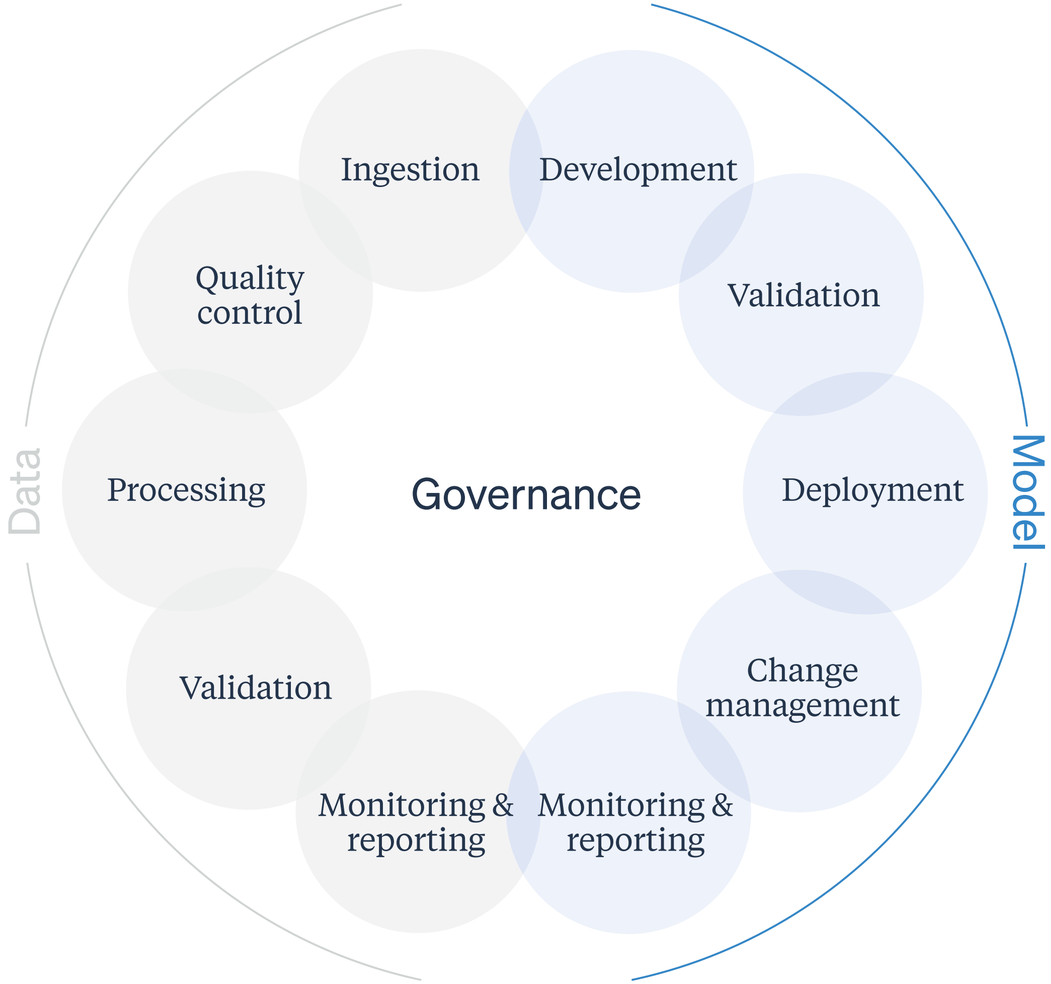

AI and data protection

AI is a top priority for the UK Information Commissioner's Office (ICO) given the high risk to individuals' data privacy rights. It has developed practical tools to help organisations ensure that AI is developed and deployed responsibly and has worked extensively with the Alan Turing Institute, in particular on "explainability", i.e. explaining how particular decisions using AI have been reached.

However, the ICO has recently warned businesses against rushing to adopt generative AI tools and solutions without considering the privacy risks, announcing "tougher checks" on whether organisations are compliant. Organisations therefore need to ensure that data protection risks are analysed and addressed at the outset and that they are asking the right (even if tricky) questions.

What are the key data protection issues in relation to how an AI tool is trained?

Data protection law is triggered not only in relation to the outputs of an AI tool, but also in relation to the personal data that has been inputted or used to train it. This creates a number of practical challenges. Have the individuals whose personal data is being used for AI training purposes been informed? Do you have a lawful basis (such as consent or legitimate interest) under the UK GDPR to use their personal data for AI training purposes? If you try to tell individuals, will this risk creating distrust and scepticism? And, if you do not have direct contact with the individuals, how will you ensure that these obligations have been met? These are difficult questions which often require detailed analysis and collaboration between all parties involved.

What are the possible data protection risks linked with the outputs generated by an AI tool?

Using AI without careful consideration could result in organisations breaching fundamental data protection principles. Data protection law requires personal data to be processed lawfully and fairly and to be kept accurate and up to date. It also protects people against having solely automated decisions made about them which have legal or significant effects.

AI systems, especially generative AI, are often trained using vast datasets, the provenance of which may be unknown. This means that important decisions about people could be made using low quality, out-of-date, inaccurate and imbalanced data. The risk here is not

AI and data protection (cont)

only receiving poor and inaccurate outputs, but also generating outputs that are discriminatory towards people based on protected characteristics (such as gender, race or age).

What are our obligations in general terms if organisations are developing AI or using or adapting AI developed by others?

These organisations will generally be "data controllers" or "joint data controllers" under data protection law. This means they will have a plethora of obligations to comply with. This includes ensuring a lawful basis, conducting data protection impact assessments, providing transparency information to individuals, mitigating security risks, complying with data subject rights, limiting unnecessary processing and ensuring that any solely automated decisions comply with Article 22 of the UK GDPR.

Where parties are joint data controllers (i.e. where the parties jointly determine the purposes and means of processing the same personal data) they will need to work together to determine who is responsible for compliance with the various obligations and document this clearly.

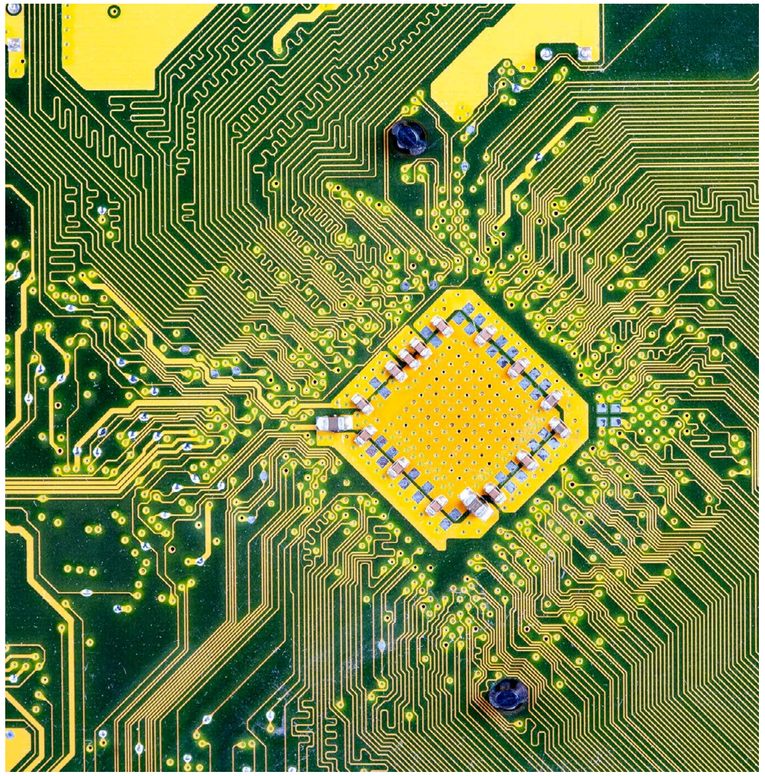

AI and contract law

The terms and conditions agreed between suppliers and customers for the development and use of AI systems to some extent reflect current drafting practice in respect of other business to business software contracts but there are AI specific aspects to consider. As the use of AI becomes more prevalent and the scope and extent to which its use becomes more specifically regulated, contract terms must evolve.

What will an AI Contract Cover?

It is essential to clearly articulate the scope of the services being provided and ensure key stakeholders understand both the technology and the benefits of procuring the AI system. The contract should make clear the type of AI technology utilised (generative or otherwise), the desired outcomes, and the responsibilities of each party involved including allocation of liability. Parties should consider including provisions related to the development, implementation, maintenance, training and support of the AI system, as relevant. Where there is a development or implementation project, clearly defining the project's goals, milestones, and deliverables will minimize ambiguities and facilitate effective performance management. It will also be important to keep an eye on EU proposals to restrict certain unfair contract terms in relation to high risk AI.

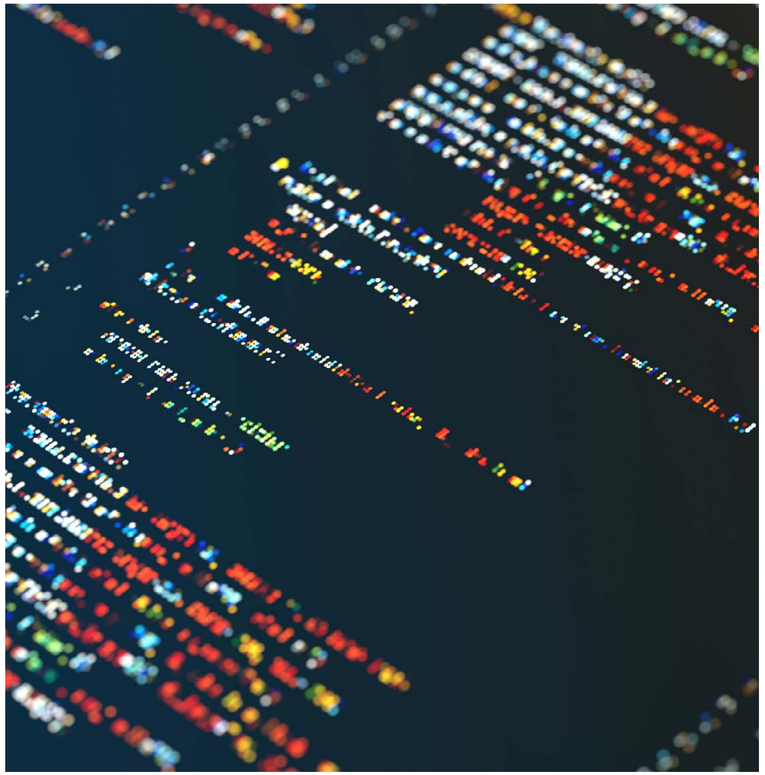

AI and contract law (cont)

Are there specific points to consider?

- **Bespoke or one-to-many models**
  Customers need to carefully consider whether the current out of the box AI models available in the market are sufficient for its intended purposes or if a bespoke product trained on a specific set of data determined by the customer is necessary. This decision will have obvious cost consequences for the customer.
- **Risks**
  It is likely that suppliers of AI technology will have to provide substantial information on their products, in particular in relation to the key principles of (i) safety, security and robustness; and (ii) transparency and explainability. Businesses procuring AI products or services will need this information to evaluate their risk. Having considered the key risks arising from the actual service or model to be deployed in the light of appropriate due diligence and consideration of the risk assessments performed by the customer, the contract should address responsibility for legal issues or liabilities arising. This may require specialist input as the cause of any liability may be complicated given the dynamic relationship between the data and AI model and who is responsible for overall governance and supervision. As such the risks of deploying an AI model in a private environment on a customer's internally held data lake will be very different from the risks of deploying a more dynamic AI model ingesting and learning from multiple public sources of data.
- **Data and IP ownership**
  The contract should specify who owns or is responsible for the data used in the AI system, rights of use for the foundation model and ownership of outputs and outline protocols for data access, storage, security, and compliance with data protection laws.
- **Open source and multi-user models**
  Users of services that interact with open source and multi-user AI models must check the terms and conditions of these models as they may require (i) warranties as to the provenance of the data input into the model; (ii) and permission to use this data for "service improvement" or for model training. This may be a concern for businesses if the data inputs include third party data that a customer may not have the right to permit the model owner to process or learn from.

Circuit breakers / kill switches

Depending on the usage of the AI system, it may be important to discuss and agree the requirements for any circuit-breaker or "kill switch" which the customer can trigger where the actions taken by the AI system risk becoming prejudicial to the customer's business. The customer should consider whether the supplier is required to maintain earlier iterations of the AI system in these circumstances. Suppliers will look to pass the costs of maintaining such earlier iterations on to their customers, as such customers will need to assess the value in maintaining these earlier iterations.

AI and contract law (cont)

What about Liability and Risk Allocation?

While traditional software agreements may focus on functionality and uptime, AI services contracts should address accuracy, training, and learning capabilities. If the AI is designed to bring about a particular result for the customer, it should consider how to define a successful outcome using measurable indicators to track performance and how this can realistically be achieved based on who is responsible for the data, model and governance.

It is unlikely suppliers will take full responsibility for their models, but they will have to commit to providing adequate information for customer risk assessments and due diligence. There will also have to be consideration of the responsibility suppliers will have to update models or to remove erroneous data. This will depend on whether a customer has decided to trust a supplier model or train its own and retain post creative processes to validate outcomes and output data.

Parties must consider potential risks such as algorithmic biases, erroneous outcomes and in the context of generative AI “hallucinations”, or the system’s inability to adapt to changing circumstances. Negotiations should include provisions that allocate risk and liability between the parties. This risk allocation will vary depending on the nature of the AI system (is it bespoke or a one-to-many service), any specific regulatory considerations, the level of spend by the customer and the customer’s expectations in terms of outcomes. Liability caps, disclaimers and specific indemnities will be focal points of negotiations, as they commonly are when negotiating traditional software agreements.

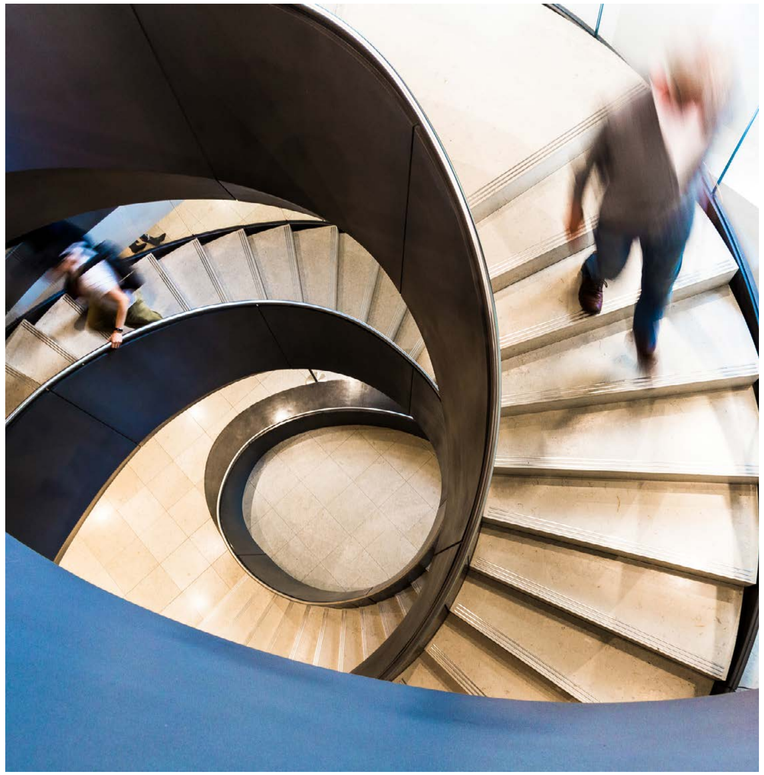

AI and intellectual property (IP)

There is nothing necessarily new about the law in the UK relating to the ownership and infringement of IP but its application to AI raises some novel and challenging issues. At the forefront of these is the question of potential infringement of IP rights, in particular copyright but also other rights related to databases, through the learning and output of AI systems.

Other important matters include the ownership of AI output and possible reforms to IP infringement exceptions intended to help put the UK at the cutting edge of AI technology.

Could the way an AI tool is trained create a risk of IP infringement?

An AI tool, if not trained on a thoroughly "clean and cleared" dataset, could certainly infringe the copyright materials used in its training. As of summer 2023, there are several ongoing legal cases or threats of litigation in the UK and US courts addressing this very issue, so additional clarification may soon emerge. While it seems to us likely that most AI training processes will involve the copying of source materials, the application of statutory exceptions to copyright infringement here in the UK, such as non-commercial text and data mining, and the temporary copies exception, will be very important. There are also other IP rights to consider such as database rights, patent rights and confidential information. The government has indicated that a code of practice will be introduced which is likely to contain proposals to facilitate the licensing of content from rightsholders for AI training. This is something that both IP owners and AI users will be keen to monitor.

What are the possible risks of IP infringement linked with the outputs generated by an AI tool?

It is certainly possible that an AI tool could produce materials that include a copy of a substantial part of an existing copyright work. Such materials will likely infringe that work. Both those who developed and train the AI and those who use the AI and provide it with prompts or instructions could theoretically be liable. A forensic approach will be required to assess the specific risks in any particular case.

AI and intellectual property (IP) (cont)

What issues arise in relation to use of third party AI tools? What part does the contract between the AI licensor and licensee play?

IP will exist in relation to the AI tool itself, for example copyright and patent rights may exist in relation to the software used. If you are using a third party AI tool, the terms of the contract are critical. First, the contract should address the issue of ownership in the software/service itself. This will usually sit with the supplier rather than the customer. Secondly, the contract should ideally also address the training of the AI: does the supplier give warranties in this respect? Thirdly, what about the outputs of the AI? The contract can override the position at law regarding IP ownership, so it is important to pay close attention. Finally, allied to these questions of ownership is that of liability. Pay attention to the liabilities of the respective parties and any limits on the same, or exclusions.

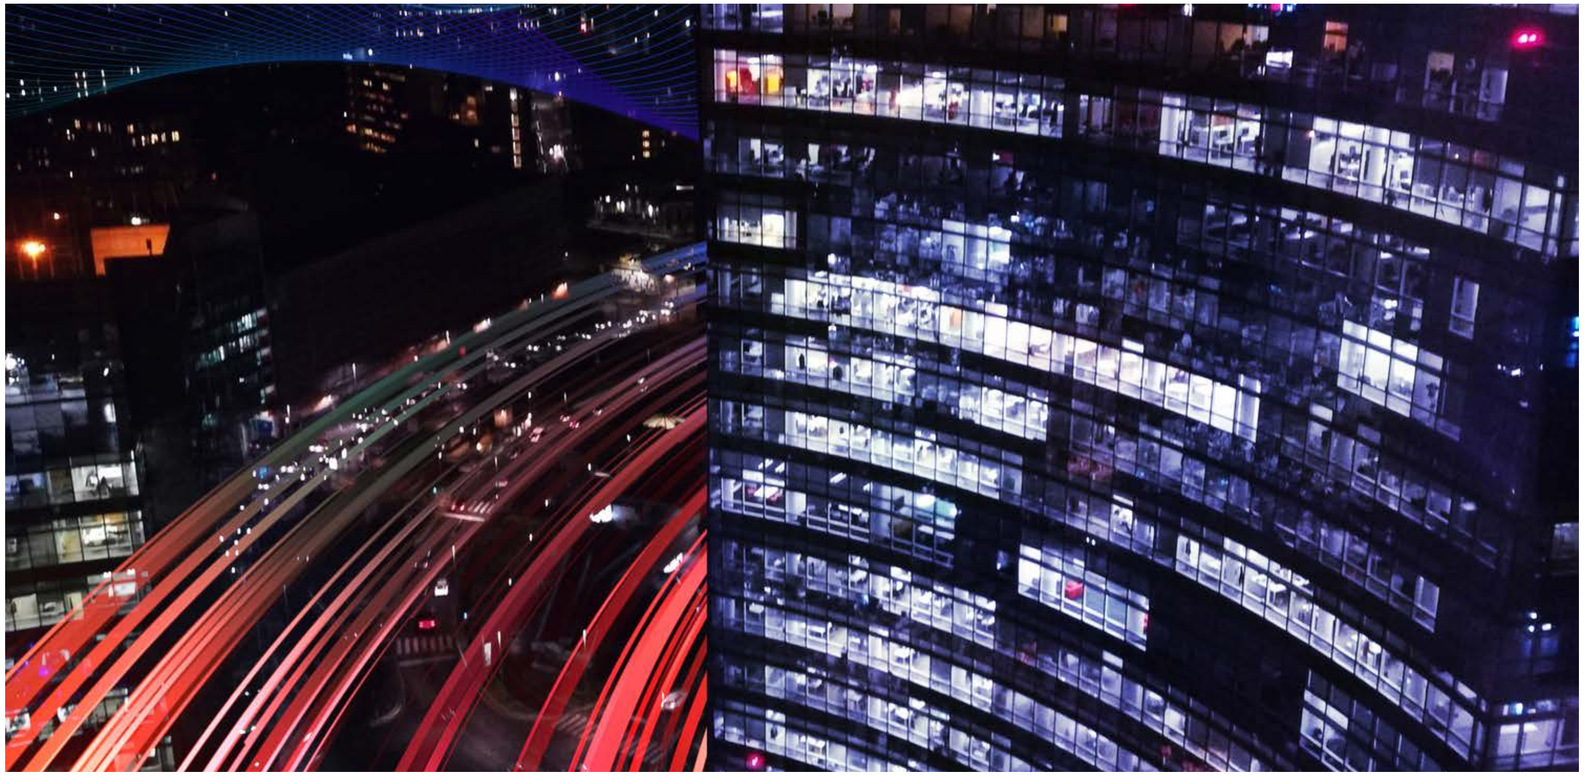

AI and employment

AI tools are used in various aspects of the employment relationship and have been for some time. From initial sifts in the recruitment process, to managing absences, automated decision making is playing a part and this will increase as further AI tools are deployed. Heavy reliance on AI tools risks undermining the personal nature of the employment relationship and the nuanced decision making sometimes required to manage a workplace empathetically.

What is the impact of AI on recruitment?

Employers are increasingly using AI tools to sift initial applications/CVs and search social media profiles for key terms. This can lead to automated decision-making where applications are rejected with no direct human involvement. As covered in more detail in the AI & data protection section, under data protection legislation there is a right (in some circumstances) to a human review of a decision that has been made by a fully automated decision-making process, but this limited right does not provide sufficient protection against potential bias in the algorithms used.

How can potential bias and discrimination in AI decisions be addressed?

Clearly AI tools are only as unbiased as the model set up and data they receive as part of their training. In a machine learning context, the potential problems were highlighted by Amazon's use of automated CV screening several years ago. Using the data from

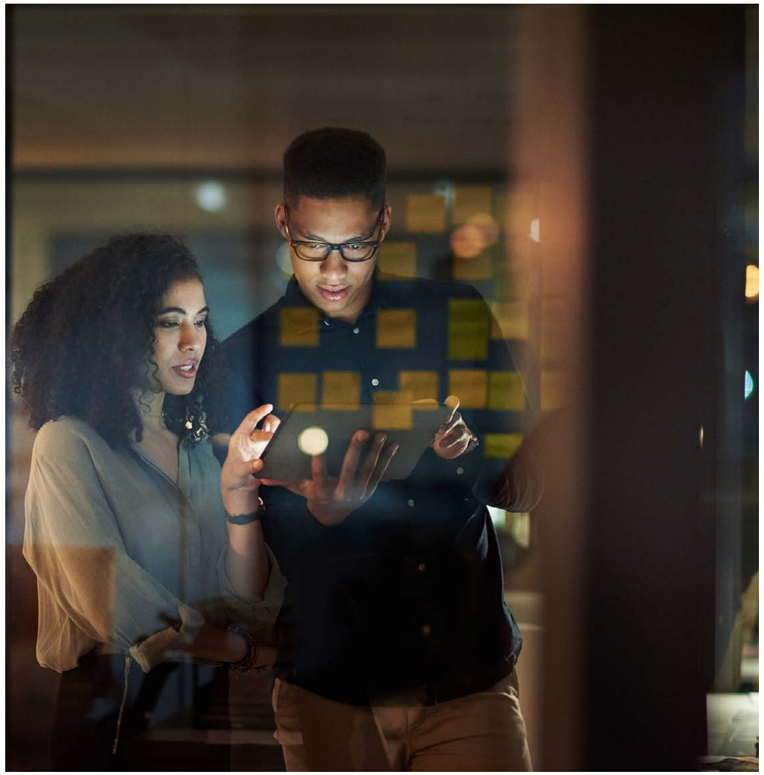

AI and employment (cont)

Amazon's historic recruitment data, the algorithm, through machine learning, "taught" itself that male candidates were preferable to female candidates. Amazon abandoned the use of the tool, but it is a warning of the potential discrimination that may arise. Where AI tools are developing as they receive information it becomes more difficult to know what the underlying algorithm is basing its decisions on, making it difficult for employers to be able to justify their decision-making process as the process becomes more opaque.

How can employers ensure compliance with employment laws and regulations?

The use of AI technology in, for example, a redundancy process, would make it much more difficult for an employee to understand if a decision to dismiss is rational and fair unless the AI model deploys appropriate transparency and explainability. The laws protecting against unfair dismissal and discrimination require an employer to act fairly and appropriately. Without a transparent and explainable understanding of the underlying model, employers will find it difficult to defend claims, leaving them exposed. From both employee relationship and risk management perspectives, employers need to be certain they can show that the decisions they make are objective and non-discriminatory.

AI and consumer law

Consumer law is based on the fundamental principles of ensuring transparency and avoiding misleading acts and omissions. The use of AI has the promise to enhance certain consumer experiences (i.e. a more personalised in-store experience) and consumer processes (i.e. more responsive complaints handling processes).

There remains however a risk that the use of AI, and in particular how it is explained to consumers, has the potential to be misleading. That, in turn, undermines the principle of transparency. As AI continues to develop, the delivery of adequate, accurate and clear information about the use and outcomes of AI will become more and more imperative.

What information should be provided to consumers?

Businesses using AI need to ensure that any terms and conditions of service provide transparent, fair and easily understandable information. Consumers must be provided with sufficient information on how their data (including personal data) will be used, whether any profiling will be undertaken on them and, crucially, what the AI itself does. It is important to keep in mind that some consumers will have limited understanding of AI itself so any explanation about AI must be easily intelligible and capable of being understood by someone with little to no understanding of the concept of AI itself. It is in addition advisable to include an explanation of the benefits of the AI to the consumer, for example, improved user experience, personalised recommendations, or enhanced product functionality. Businesses should also specify how the terms may be updated or modified in the future in the event the nature of the AI itself develops. Terms should include a mechanism to notify consumers about significant changes and allow them to review and accept the revised terms.

AI and consumer law (cont)

What other issues should be considered?

- Privacy and data protection
  As noted in the AI & data protection section, the use of AI must be in line with applicable data protection laws. Adopting a privacy by design approach will be crucial for such a rapidly evolving technology and will ensure that businesses can effectively provide the necessary updates to consumers. This is particularly important where businesses link personal data to send targeted marketing materials.
- Discrimination and bias
  Consumer laws prohibit unfair and discriminatory practices and AI systems have in the past been known to discriminate against consumers based on protected characteristics such as race, gender, or age. Businesses need to ensure that they can adequately explain how AI-based decisions are made and provide avenues for consumers to seek clarification or challenge decisions that significantly impact them. Use of AI should always be underpinned by a commitment to fairness and equal treatment for all consumers. See the AI & employment section for more on the issue of bias.
- Product liability
  There are questions over whether existing consumer laws address questions of product liability adequately. No new regulation has yet been proposed in the UK but in contrast the EU has introduced a new AI Liability Directive. This is an area that will come under increasing scrutiny as the use of AI continues to grow but will be one of the primary routes for redress.
- Price discrimination
  Again, consumer laws prohibit unfair pricing practices and price discrimination. Price discrimination is not a new concept, but AI can potentially enhance its effectiveness and efficiency. By way of example, AI can use individual customer data, such as browsing history, purchase history, demographics, or even social media activity, to tailor pricing specifically to each consumer to identify, and charge them, the highest price they are willing to pay while still making a purchase. However, the development of this algorithm must not be based on protected characteristics such as race, gender or age which may perpetuate existing inequalities.

AI in mergers & acquisitions

AI has been an important and increasingly prevalent feature of M&A transactions for some time. Not only will a buyer need to understand how AI is used in the target business (to understand and evaluate potential risks arising from that use and to plan for how the AI assets of the target can best be integrated with those of the buyer), but buyers and sellers are increasingly using AI systems in various aspects of the transaction itself.

How is AI used to implement an M&A deal?

M&A lawyers have in recent years been deploying AI tools (particularly machine learning technologies) to process large amounts of due diligence (DD) data in order to identify issues that are the focus of their DD processes. There are a number of commercially available products that will read the entire contents of a data room and more or less immediately identify features which would otherwise take individuals significant amounts of time to find. Until recently it has generally been machine learning that has been used in this kind of review, but generative AI solutions are emerging in this area as well.

The technology can also quickly identify clusters of documents that are in a similar form (for example they might all be based on a template distribution agreement) and then equally quickly highlight variances – for example documents which are missing wording (or even whole clauses), documents which have additional clauses and documents where template clauses have been varied.

Similarly, the technology can screen a data room to pull out all of the documents which contain specific types of clause, such as a change of control clause. This requires the technology to utilise learning that can do this (as a change of control clause does not need to contain the word “change” or “control”). The model is trained on document sets that have gone before and if it is in doubt it checks with the human user – based on the response it is given it “learns” by refining itself, having taken into account the new information.

AI solutions can also assist in the preparation of draft deal documentation and to overlay existing machine learning tools in the comparison of draft documents (received from the other side of the transaction or produced internally) with a specific data set. This could for example be the internal library of similar documents held by a law firm or legal department, or an external data set such as documents filed with the SEC in the US and stored on its Edgar system. Clearly defined policies and governance will be required in

AI in mergers & acquisitions (cont)

order to deploy this technology and explain to clients the benefits and risks involved.

How can AI-related risks be identified and mitigated?

A key component of any DD review going forward will be the assessment of AI-related risks in a target. For example, are the target or its personnel deploying AI in a way that breaches the law or the rights of third parties? There are many examples in this guide of how such violations might occur and of course the particular issues will often depend on the type of the underlying business. Consider the following:

- Is AI being used in a business in a way that is potentially discriminatory (for example in making decisions between job candidates or, in financial services businesses, when deciding whether to give a person credit)?
- Are AI products being used by target personnel in a way which might compromise the confidentiality of confidential information, or which might enable AI platforms to use that data in potentially unexpected ways?
- Is AI being used in a way that unfairly processes personal information in breach of data privacy requirements?
- Do the issues around the potential inaccuracies produced by generative AI ("hallucinations" etc) create risks for the business?
- Does the use of AI in the business risk infringing the IP rights of third parties, or of itself creating IP which is not sufficiently protected?

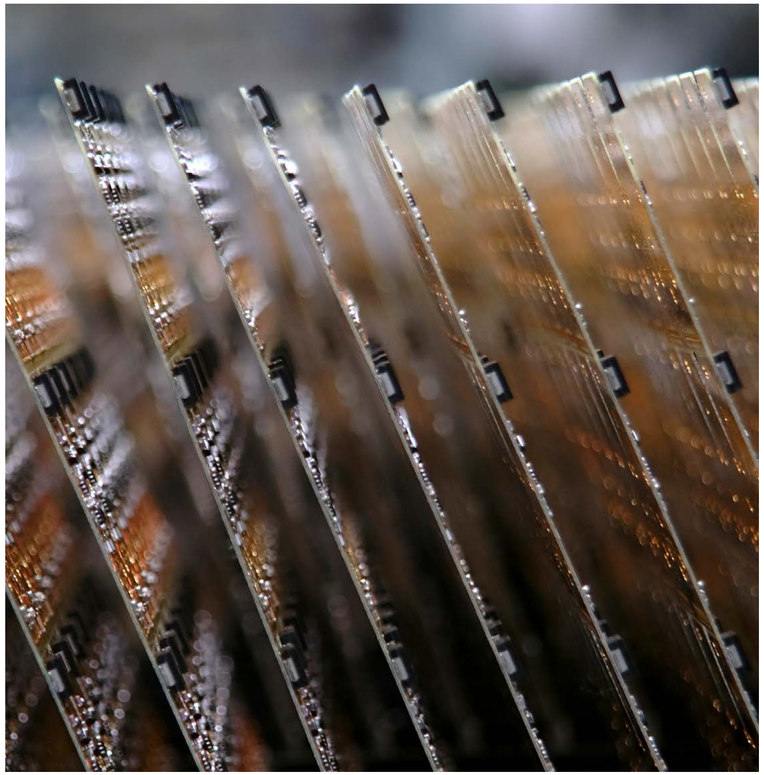

AI and dispute resolution

Technology solutions have been used for some time to streamline and assist litigation processes, from Technology-Assisted Review for disclosure exercises, to expert selection analytics tools, compensation calculators, and case law analytics tools. But as the power and sophistication of the most advanced AI tools continues to expand, how might its impact on the world of dispute resolution increase in the future?

Is the judge of the future artificially intelligent?

In some cases, the judge of the present is artificially intelligent. In Europe, Estonia is at the forefront of this advancement and has recently implemented a system which allows AI to issue a decision in cases involving disputes valued below €7,000. In China, the Hangzhou Internet Court uses virtual judges to reach decisions in disputes involving digital matters. This intervention is usually in lower value cases and the typical caveat is that the AI decisions are appealable to a human judge. Indeed, Sir Geoffrey Vos, Master of the Rolls in the English Courts, recently commented that "AI will be used within digital justice systems and may, at some stage, be used to take some – at first, very minor – decisions... The controls that will be required are (a) for the parties to know what decisions are taken by judges and what by machines, and (b) for there always to be the option of an appeal to a human judge." Sir Geoffrey also expressed the opinion that, despite attempts to advance the systems in place, the processes in the Business and Property Courts are no

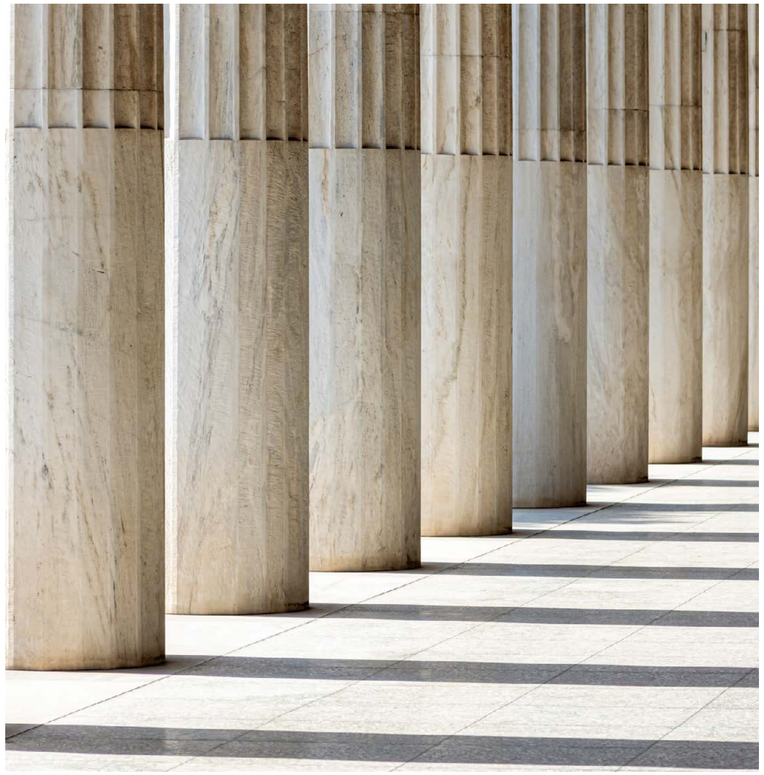

AI and dispute resolution (cont)

quicker than 20 years ago, and that "we need to re-think the process for the digital age".

If the use of AI is inevitable, what are the risks and how can we mitigate/protect against those risks? There are a number of risks associated with the use of AI in dispute resolution. One example is the well-publicised concerns over discriminatory decision-making by AI, where bias in the data on which the AI relies is reproduced in its output. A more recent source of concern is AI's apparent ability to embellish – with one lawyer in the US inadvertently citing fake case law generated by ChatGPT, which had assured him the citations were genuine. The key precaution is to fully understand the parameters and limitations of the software being used, and to ensure that AI-generated results are carefully cross-checked against more tested resources.

Responsible governance is also critical. In the UK, the government White Paper, referred to in the AI regulation & ethics section, proposed that AI will have neither specific regulation nor a specific regulator, but rather be monitored by sector-specific regulators. This suggests that rules on the use of AI in dispute resolution may be a matter for professional regulation and court-approved processes.

Who is liable for AI-produced errors?

Liability for AI errors will of course depend on context, but the tort of negligence appears likely to be a typical cause of action, unless and until statutory measures are introduced. In essence, the allegation in negligence would either be that:

1. the developers of the AI owed a duty to the user and failed to take proper care that the AI would be reliable, or
2. where the issue is a service/product supplier relying on AI, that the supplier was careless in relying on AI to perform its role.

It is worth noting, though, that the European Commission has considered this issue at length (including whether to introduce strict liability offences for faulty AI) and has proposed legislation – so statutory intervention may well be forthcoming in many jurisdictions.

Charles Russell Speechlys

Contact us

Mark Bailey
Partner

Mark.Bailey@crsblaw.com
+44 (0)20 7427 6519

Janine Regan
Legal Director

Janine.Regan@crsblaw.com
+44 (0)20 7427 1074

Nick White
Partner

Nick.White@crsblaw.com
+44 (0)20 7438 2294

1213106641_V2