



Informing Progress - Shaping the Future

FOIL Update 7th August 2025



AI and the Law: The Law Commission Discussion Paper

The Law Commission's discussion paper, *AI and the Law*, represents a significant step in framing how the legal system of England and Wales should engage with the growing influence of artificial intelligence (AI). While not proposing formal reforms, the paper outlines the legal questions AI raises and highlights the need for clarity, accountability and regulatory vigilance.

With AI systems becoming increasingly autonomous and capable of influencing or making decisions independently, the paper underscores the legal system's growing need to evolve with technological developments to ensure it can mitigate the risks while embracing the benefits.

The Rise of AI

The paper opens by drawing attention to the impressive capabilities of AI, from the deployment of autonomous vehicles to advances in medical diagnostics and the rapid proliferation of large language models (LLMs). In delivering these, AI has driven significant improvements in productivity, innovation and investment across multiple sectors, with notable examples being the widespread adoption of LLMs such as OpenAI's ChatGPT or Google's Gemini. The expectation is that the rapid development of AI will continue to fuel considerable social and economic benefits as its acceptance grows.

However, alongside these benefits comes the risk that AI will harm society, with AI implicated in incidents involving fraud, cyberattacks, disinformation and discrimination. The environmental impact of high-energy AI infrastructure, the threat of workforce displacement and the potential for autonomous systems to act in unpredictable or even malicious ways are all raised as legitimate areas of concern. These developments have prompted regulatory responses globally, including the European Union's AI Act and its accompanying Code of Practice for general-purpose AI models.

In this context, the Law Commission calls for a broader legal conversation around how existing legal doctrines and frameworks can accommodate a new class of non-human actors that are capable of learning, adapting and making decisions.

Autonomy, Adaptiveness and the Issue of Liability

One of the central themes the paper explores is the growing autonomy and adaptiveness of AI systems. AI is no longer limited to executing pre-programmed commands, but can learn from data, develop novel strategies and operate with minimal or no human oversight. This autonomy complicates traditional legal reasoning around accountability; for instance, if an autonomous AI system causes harm, does liability rest with the developer, the deployer, the user or none of the above?

The report raises the danger of 'liability gaps', or situations where no natural or legal person can be held accountable for harm caused by an AI system. These gaps could undermine public confidence, impair access to compensation and disincentivise responsible innovation.

While parallels are drawn to how corporations, as non-natural persons, have been integrated into legal doctrine and are held responsible for their conduct, current AI systems lack legal personality and cannot themselves be held accountable. The Law Commission suggests that this issue is likely to occur more as AI capabilities evolve, raising the acute possibility of specific AI systems being granted legal personality in future.

Complex Supply Chains and Causation

AI development is rarely the product of a single entity, but involves a complex supply chain of data collectors, model developers, software integrators and end-users, each possibly based in different jurisdictions. This fragmented landscape presents significant complications when attributing legal responsibility.

For example, if a healthcare provider uses an AI diagnostic tool that delivers a faulty output that leads to patient harm, establishing liability can prove problematic. With a provider, software company, model developer and perhaps a separate training data supplier, deciding who should bear liability for the harm caused is not straightforward.

The common law principles governing duty of care, proximity and foreseeability must be reevaluated in the light of these multilayered relationships. While the complications caused

by complex supply chains are not unique to AI, they illustrate how decisions upstream in the chain can have severe consequences for those involved downstream.

Establishing causation poses further challenges, as AI systems cannot be held liable, meaning the focus shifts to the person liable for that system. In negligence claims, claimants must prove that harm would not have occurred “*but for*” the defendant’s actions and that such harm was reasonably foreseeable. The complexity and unpredictability of AI outputs complicate this analysis, as AI systems can behave in unexpected ways, meaning that even models tested extensively may cause unforeseen harm. The paper stresses that while not all unexpected harms necessarily result in liability, legal certainty is key to ensuring access to remedies and enabling insurance coverage.

Mens Rea, Recklessness and Accountability

The paper goes on to examine how AI affects private and criminal law with respect to offences requiring a mental element, such as intention, knowledge or recklessness. As AI systems do not possess consciousness or intent, criminal liability must instead be attributed to a natural person or corporate body.

This creates issues with autonomous and adaptive AI systems that have no human input. The paper highlights an example where an AI system makes misleading financial statements that breach criminal law. In this scenario, the lack of human input would make it difficult to establish that a firm knew a false statement was made or that any individual within the organisation had the requisite knowledge or acted recklessly.

These complexities in establishing recklessness, causation and liability highlight the potential need for new legal tests or statutory duties to ensure accountability in the age of AI automation. As highlighted above, the intricate steps involved in the supply chain, some of which may be repeated as AI systems are deployed, make the task of determining who is legally responsible extremely challenging.

The Problem of Opacity

Many AI models, particularly those based on deep learning and large-scale data, produce outputs through highly complex and often impenetrable processes. This lack of transparency, whether due to commercial confidentiality or technical complexity, has profound implications for justice and accountability.

In public law, for example, the use of AI systems in administrative decision-making can impede judicial review. It may be impossible, for example, to determine whether a public body made a lawful decision by considering all relevant factors and avoiding irrelevant ones. Such decisions can be questioned, regardless of whether automated AI was used or whether a human decision was reached with input from AI. Either way, it undermines the transparency and procedural fairness fundamental to lawful public decision-making.

Concerns also surface in criminal trials, particularly around the admissibility and probative value of evidence generated by AI systems. If an AI-assisted tool informs a sentencing decision or influences a prosecution, defendants must be able to challenge its basis. Without logical reasons or accessible data, this right may be meaningless and jeopardise the principle of a fair trial.

Oversight, Reliance and Professional Duties

The paper makes the distinction between scenarios where humans retain oversight of AI systems and those where automated AI is employed with little or no human intervention. In both cases, the challenge lies in determining the appropriate level of human responsibility based on the scope and content of their duties.

In some cases, it is straightforward; for example, lawyers are expected to verify the accuracy of legal authorities cited in court submissions. However, there have been instances of fictitious citations generated by AI tools presented in court. These have led to disciplinary action and emphasise the risks of delegating professional duties to machines.

Other situations, however, are less clear. A medical professional using an AI tool for diagnosis might reasonably rely on its recommendations for subsequent treatment. Yet, if the AI fails and results in harm, should the practitioner be liable for not exercising their independent judgment and relying solely on output from the system? What if the situation were reversed, and they used their own judgment, particularly where there is evidence that AI analysis is demonstrably more accurate than that of a human? The paper highlights that professional and regulatory guidance may be needed to provide clarity in these situations.

Data, Bias and Discrimination

A further legal challenge arises from the data used to train AI systems and the potential for issues related to data protection and copyright. AI models are trained using vast datasets, and these may include copyrighted material or personal data, raising issues under copyright law and UK GDPR. Informed consent, lawful processing and data minimisation are all principles that may be difficult to comply with when dealing with the opacity and scale of the training data used by AI models.

Bias in training data is also a persistent concern, with AI systems capable of reproducing and sometimes amplifying societal biases. This can lead to discriminatory outcomes in recruitment, policing, healthcare and public service delivery. Even when unintentional, such outcomes may contravene equality duties and lead to unlawful discrimination.

Towards a Legal Framework Fit for AI

The Law Commission discussion paper does not offer all the answers and stops short of proposing specific reforms. It does, however, identify several issues by AI and raises broad questions about how legal duties should be determined. Furthermore, it sets out potential routes for legal development, including:

- Reviewing civil and criminal doctrines of liability to ensure they can accommodate autonomous and adaptive systems.
- Exploring whether certain categories of AI systems should be granted some form of legal personality.
- Developing clearer professional and regulatory standards on oversight, risk assessment and reliance on AI tools.
- Enhancing transparency and accountability mechanisms in public decision-making where AI is used.
- Addressing issues of bias and discrimination through proactive design, testing, and monitoring of AI systems.

As AI systems continue to evolve in capability, complexity and societal reach, the legal frameworks that govern them must evolve to keep pace. Multi-stakeholder engagement, involving government, industry, academia and broader society, is needed to address the legal uncertainties and structural risks posed by AI to affirm the sector's role in safeguarding rights, ensuring accountability and fostering responsible innovation

This publication is intended to provide general guidance only. It is not intended to constitute a definitive or complete statement of the law on any subject and may not reflect recent legal developments. This publication does not constitute legal or professional advice (such as would be given by a solicitors' firm or barrister in private practice) and is not to be used in providing the same. Whilst efforts have been made to ensure that the information in this publication is accurate, all liability (including liability for negligence) for any loss and or damage howsoever arising from the use of this publication or the guidance contained therein, is excluded to the fullest extent permitted by law.