

November 2020

UK Intellectual Property Office – call for views on artificial intelligence: Comments provided by Mills & Reeve

INTRODUCTION

This document is in response to the *Artificial intelligence and intellectual property: call for views* issued by the Intellectual Property Office in September 2020 (the “**Consultation**”).

Mills & Reeve is a national UK law firm with 133 partners and a total strength of over 1,000 staff operating from six offices including London, Manchester, Birmingham and Cambridge. Mills & Reeve is one of the top performing law firms in the UK when it comes to client satisfaction, according to the latest editions of legal directories Chambers UK and The Legal 500, and has been named for a record seventeenth year running as one of the 100 Best Companies to Work For in The Sunday Times annual survey. Mills & Reeve acts for a range of clients who have an interest in the development of artificial intelligence (“**AI**”). We advise a range of clients on issues relating to AI, from automated transport to healthtech, and therefore have a close interest in seeing that a robust legal and regulatory framework is put in place.

We have carried out a series of discussions with clients and contacts to identify issues and concerns for inclusion in our submission. Many organisations that would be expected to be closely interested in developments in this area are not making their own submissions. This is partly because intellectual property law is seen as technical and complex and is not readily understood by innovators in the field. It also stems from a widespread view, which we will discuss further below, that traditional intellectual property is not a good fit for AI innovation and not something that is of great interest to developers.

We begin with general comments on the approach taken. Our comments are confined largely to the use of the IP system for protecting technical innovations, rather than artistic creativity. We therefore focus on patents, copyright and related rights, designs and trade secrets, and will not address the trade marks section. The Consultation questions are included in italics for ease of reference.

GENERAL COMMENTS

AI innovators face a unique set of issues when protecting their data, technology and inventions. This is a fast-moving area – innovations may only have a short effective lifespan. Ideally, innovators want to protect their “patch” so that competitors cannot benefit from the same approach. In theory, patents are ideal for this purpose as they offer strong monopoly protection for inventions. However, the slow-moving and costly patent process is designed with a well-defined invention and a 20-year time scale in mind. In addition, the current exceptions to patentability mean that many novel applications using AI cannot easily be protected. Patents are not, therefore, always a good fit for AI and machine learning innovations. Copyright and database right offer some protection for software and databases, but they leave many gaps and areas of uncertainty. This is a good moment to assess whether the existing IP system works well in supporting AI innovation, and consider what improvements could be made.

As a general comment, we would emphasise a lack of engagement with development of intellectual property law in this area amongst those organisations most likely to be affected by it. The system is seen as complex and difficult to navigate. Patent protection is not often seen as a good fit, primarily due to the speed of development in the sector (although we note that developers in other parts of the world, particularly China, are very active patent filers in the AI arena). It is also seen as complicated and costly to use, although it is relied on for selected, significant developments.

Copyright is relied on for software protection, but the sui generis database right is little known and indeed, difficult to apply successfully. Data assets are of key importance in the AI field, but are not well protected. As a result, reliance on trade secrets and confidentiality agreements is widespread. Although this does offer an effective form of protection for many innovators, it has serious drawbacks. Innovative AI developments, and important datasets, are not made widely available but are closely guarded for use by one or few operators. Further, disclosure of key material, whether fraudulently, inadvertently or in response to a legal obligation, can effectively remove all protection.

There seems to be an underlying view that AI innovations can fit into the IP system or the IP system can be made to fit around AI innovations. However, this is only true in some situations. For example, patents may be available where technical innovations result from the use of AI, such as novel drugs/drug applications, or control of technical processes based on ability to analyse based on machine learning.

As we have said, many do not see the current system as providing a good match for broader business or innovator needs, either in terms of the requirements for obtaining protection, the type of protection obtained, its enforceability, and the duration of that protection. For example:

- a) there are significant questions as to whether protection should be available for AI innovations with outputs in the social/behavioural field; and
- b) there are issues around excluded subject matter, whether the term of protection would be too long, enforceability and indeed whether such protection promotes or inhibits innovation in relation to innovations in the internal operation of AI systems.

As a result, many innovators, and in particular smaller businesses see little or no benefit in formal protection such as patenting, do not give it consideration or feel that it is not really fit for purpose in relation to AI innovations. Most contacts that we spoke to said that they would rely on confidentiality rather than the patent system. Some contacts were worried that the patent system would be open to abuse (such as aggressive litigation by patent assertion entities, or “trolling”, or attempts to dominate platform innovations at the expense of less well-funded market participants).

Important playing pieces in the innovation/competitive battle are seen to be:

- confidentiality;
- the control of data and its counterpart, access to data;
- contractual frameworks for access and reward (and limitations on the enforceability of contractual frameworks);
- access to skilled people;
- speed to innovation and to market; and
- market power.

There might be scope for development of a specific right to protect some forms of innovation in the AI field. A key drawback of this approach would be its application to the UK only (at least for the present) whereas

AI is unquestionably international. And there are real concerns that protection in this field may inhibit innovation or limit its wider application, rather than promoting it.

While the current system is seen as ineffective in many respects, we would suggest caution in the development and implementation of any new rights. This would require a careful analysis of how these might incentivise particular behaviours, and influence the desirability of the UK as a location for development and operation of AI technology. The problems created by too many rights, rather than too few, can seriously impede smaller innovators. Arguably, competition to be first to market may be a better incentive to innovate.

DETAILED RESPONSES

Patents

Encouraging innovation involves two, sometimes contradictory, elements. It is of course important to provide innovators with incentives for their investment of time and money, and their inventive spark. This means offering an intellectual property system that protects valuable developments with an appropriate degree of exclusivity and recompense. But equally important is the need to prevent overreach of these rights to the detriment of other developers. Any intellectual property system must recognise these two priorities, and aim to strike a good balance between them.

A concern often raised by smaller innovators is the “patent thicket”. Poor quality patents, granted too easily, can act as a block on innovation. We consider that an overgenerous approach to granting monopolies, whether patents or an alternative, new right tailored for this field of innovation, risks acting as an obstacle to real innovators. Smaller innovators can be constantly beset by demands for licensing royalties, and often lack the financial resources to engage in infringement and validity litigation, and so can be pressured into reaching a commercial settlement and paying royalties in respect of patents or other intellectual property rights that would, if tested, be found invalid. Portfolios of patent rights are often bought up by patent assertion entities, whose business model involves use of litigation or the threat of it to derive revenue from genuine innovators. We consider this to be a practical unfairness in the system, best addressed by avoiding overgenerous granting of monopolies, and vigilant surveillance of potentially anticompetitive practices. Rigorous patent examination is a key element of this.

The UK has been very successful at attracting AI-rich industries such as Fintech, internet search and voice recognition. One school of thought sees an approach that develops and promotes innovations rapidly, and acts fast to establish them as first mover to see them established as the market leader, as likely to be more successful than reliance on IP protection. Patents should be available, but for substantial technical advances only.

Of course, measures taken by the UK alone are unlikely to influence the development of the industry globally. As we have seen, patent litigation in the mobile telephony sector is increasingly international, with global royalty setting now the norm. However, providing an environment for innovators to flourish without unnecessary obstacles is likely to best promote healthy development of the sector.

1. What role can/does the patent system play in encouraging the development and use of AI technologies?

The Consultation highlights the aim of ensuring that the UK has the best environment for developing and using AI. While strong patent protection can offer a valuable incentive for innovators, it is equally

important to avoid granting excessive protection for inventions. Progress depends on innovators building on existing technology to make new advances. Granting 20-year monopoly protection in a fast-moving field is likely to stifle innovation rather than promote it if the protection is broader than is merited by the disclosed invention.

In evaluating whether the patent system plays a role in encouraging the development and use of AI technologies, it is helpful to distinguish between different sorts of innovation arising from AI, and then to look at the creative, investment and business environment relevant to those types of innovation.

Innovations arising from AI may be classified as follows:

1. AI used to identify a problem, to identify products, processes or services, or which achieve a solution to a particular problem, in each case external to the AI itself;
2. the development of the “AI engine” whether to make it:
 - 2.1. more efficient, for example, able to process larger amounts of data more quickly and so identify patterns otherwise difficult to identify; or
 - 2.2. able to provide “better solutions”, such as improved standards of image recognition, discrimination, prediction or emulation; and
3. innovations in the system inputs, such as insights into the input data used for training an AI engine.

Class 1 – AI as a tool

With class 1, the value chain will begin with the product or service. Here the AI is a tool and value is generated by the product or process. The AI system itself may or may not be innovative. The end-product innovation may instead be a result of the data used, or the knowledge and skills of the developers. Existing IP protection can be employed in a conventional way to protect this class of innovations, with the AI system simply a tool in the development process. However, some innovations of this type may not be protectable under the existing patent regime, notably when they are essentially business methods or similar “non-technical” applications. In that case, it will be necessary either to look for an innovation in some other aspect. Alternatively, it may be considered appropriate to provide some other form of protection or extend the scope of patent protection in order to incentivise certain kinds of innovations e.g. crowd behaviour simulations, or direct diagnosis tools.

Class 2 – development of the “AI engine”

Here, different considerations apply. These might be “platform” innovations – useful in a wider field than the developer’s area of focus, or useful where a developer focuses on providing services to entities in other fields. Given the current fast pace of development, many of these innovations are likely to be short-lived. It may also be difficult to determine whether others are using them. As a result, striking the right balance for the optimum level of protection changes. Providing enough enforceable protection to incentivise innovation, while avoiding anti-competitive behaviour becomes more difficult.

Class 3 – system inputs

Class 3 faces similar problems to class 2, but with greater difficulties in finding suitable forms of protection. In many situations, class 3 innovations are seen as the area of greatest innovator value: possession and control of the data used.

A lack of interest?

Given the difficulties with protecting each class of innovation, it is perhaps not surprising that the feedback from our discussions with those active in the AI field is that the current patent system does not provide them with a meaningful incentive for innovation. “Lack of interest” characterised many responses. In more detailed discussions, a large proportion of our contacts indicated that they would rely on confidentiality to protect their innovations in AI. Protection for products or services which AI had been instrumental in “discovering” or characterising would be sought in the usual way.

2. *Can current AI systems devise inventions? Particularly:*
 - a. *to what extent is AI a tool for human inventors to use?*
 - b. *could the AI developer, the user of the AI, or the person who constructs the datasets on which AI is trained, claim inventorship?*
 - c. *are there situations when a human inventor cannot be identified?*

Based on our discussions and experience with businesses our responses to these questions are:

- a) Yes. AI is used as a tool for human innovators.
- b) Any of these (or more than one of them) could be the deviser of an invention, depending on their involvement, and if that is the basis for claiming inventorship, could so claim.
- c) Yes, in the sense that it may not be possible to identify who “devised” an invention. However, as explained in the more detailed discussion, that is probably already the case with some inventions currently outside the field of AI. We consider it likely that the courts dealing with disputes will nevertheless be able to decide who should be treated as the inventor. Further, in most jurisdictions this will not matter as long as underlying contractual arrangements deal with ownership of any inventions.

3. *Should patent law allow AI to be identified as the sole or joint inventor?*

Not in so far as this may give rise to any personality rights for the AI (discussed further below). If this is a short circuit to providing a base for determining ownership of the rights in the invention, it may be convenient. However, as explained in relation to ownership our view is that this will not necessarily confer rights on the most appropriate entity.

Many currently predict that AI systems with greater autonomy (often referred to as artificial general intelligence, and at a step further on, artificial superintelligence) will be developed. If this happens, the inventor question becomes more difficult. If an AI systems is able to “think” as a human does, and apply its intelligence to any problem, then its human developers can no longer realistically be deemed to be inventors of what it may produce.

Simply disallowing patents where no human inventor can be identified risks reducing the incentive for innovation, and for disclosing the innovation to the public. Some suggest that the AI itself should be granted the status of “inventor”. We do not see this as the right course. It could lead to a situation where the AI is entitled to own the patent and is granted an inappropriate level of equality with humans. Indeed, we question whether an innovation that has been developed by an AI system where the input from its human developers is not inventive should be classed as an “invention” at all. Where the algorithm, the selected training data and the problem being addressed are obvious to try, with any innovative elements being generated by the system, this should not be classed as an invention capable of patent protection.

What is an invention?

Stepping back, it is worth considering what an invention actually is. If there is no human involvement, can we say that an invention has occurred at all?

This may be imbuing the concept of an invention with more substance than, in reality, exists. In the pharmaceutical sector, for example, an invention may be made by carrying out mass throughput screening for chemical compounds with particular characteristics.

Indeed, the concept of “devising” an invention, and who is the “deviser”, is in many ways already outdated. For example, an invention may arise from an observer seeing an unusual event (as in the discovery of penicillin), or by carrying out systematically a large number of experiments, one of which turns out to be successful. An innovation may arise from ideas put forward by a number of individuals over a period of time gradually crystallising into a coherent idea. Nevertheless, it is useful as providing a (theoretically) identifiable construct on which to base claims of ownership.

Who is the inventor?

Who such a person might be depends on the type of AI invention contemplated.

Where the innovation relates to the internal operation of the AI engine (class 2, above), then the deviser will likely be the creator or improver of the AI.

For class 1 innovations – where the AI produces an output which is an innovation – the deviser could be understood as the person(s) who put in place the arrangements giving rise to the innovation with some foresight as to the characteristics of the output, or selectivity applied to/appreciation of the output when it arises. We already see ever more sophisticated tools used in generating innovations.

In this case, it is most likely to be the user of the AI, if that person controls its operation. It could also be the provider of the datasets, if the operation of the AI is purely mechanistic when applied to the datasets. There is some risk of uncertainty in this context, particularly in relation to artificial general intelligence.

4. *If AI cannot be credited as inventor, will this discourage future inventions being protected by patents? Would this impact on innovation developed using AI? Would there be an impact if inventions were kept confidential rather than made public through the patent system?*

In terms of promoting innovation, the ideal foundation is to credit either:

- a) the person most closely connected with the requirement to invest in order to generate, develop and exploit an innovation; or
- b) the person most likely to have ideas which will lead to innovation.

However, each of these might lead to uncertainty and who those persons are will depend on the field of innovation.

For practical reasons, and based on relatively weak observations from users, the inventor/ owner is most usefully identified as the person conceiving or putting into practice the innovation. In the absence of such a person being identifiable, as may be the case with artificial general intelligence, it should be the person making arrangements which resulted in the innovation being made. It may be desirable to give specific guidance or “deeming” provisions towards this conclusion.

An alternative approach would be a “reward for investment” model instead of patent protection. This could be modelled on the EU system of protection for databases. Ownership of rights to work product would pass to the person or organisation that developed the AI system. This is a more logical step than conferring inventor status on the AI system itself.

5. Is there a moral case for recognising AI as an inventor in a patent?

In short, no.

There is not a strong moral case for recognising an AI as an inventor. We are not aware of a substantial body of opinion calling for AI systems to be afforded moral rights and we do not consider this to be a useful approach. The rationale for recognising a human inventor, whose future career and employment prospects are likely to be enhanced by the status of inventorship, does not hold true in the same way for AI systems. Granting inventor status offers a degree of personhood to the AI system that is not consistent with its nature as a product of human work and ingenuity. Neither can we see a compelling legal or procedural reason for recognising an AI as inventor.

Further, the idea of conferring moral rights has a corollary of expecting that AI (agents) will assume moral obligations. Doing so runs a significant risk that human developers will allow themselves to shed moral responsibility in favour of the AI, and enable or permit behaviour which the developers should have ensured is compatible with human moral codes.

6. If AI was named as sole or joint inventor of a patented invention, who or what should be entitled to own the patent?

As discussed above, we do not consider it to be appropriate for an AI to be named as an inventor. However, if this course were to be adopted, then ownership of any resulting patents should lie with the individual or organisation primarily responsible for the AI’s activity which resulted in the invention.

In complex situations, there may be difficulties with determining which person or entity this is. There could be substantial contributions from, say, the creator of a database, the designer of an AI system and the person who decides to apply the AI system to the database. Additional rules might be needed to define these requirements. Indeed, it would be possible to have joint ownership where there have been several different contributions to the AI’s activity. However, this is not a new concept and should not lead to insurmountable difficulties. Patent law currently recognises joint inventorship and co-ownership where different individuals are involved.

7. Does current law or practice cause problems for the grant of patents for AI inventions in the UK?

8. Could there be patentability issues in the future as AI technology develops?

The initial requirement for obtaining a patent is to file an application and to name the inventor. If the inventor cannot be determined, then this does create a problem. As noted above, however, we believe the courts will be constructive in working out a solution to who the deviser of the invention is.

In most jurisdictions, it is only another person with a claim to be the inventor or entitled to the rights arising from the invention who can challenge the naming of any individual(s) as inventor. Therefore, assuming the applicant is willing to identify some individual(s) as inventor, this should not present a problem. (The position is different in the United States, where naming the wrong inventor can be a basis for invalidity. There is therefore some need to resolve this on an international basis.)

In addition, two of the basic requirements of patent law present particular difficulties in relation to AI inventions. These are:

- the need for an inventive step; and
- the need for an enabling disclosure.

The **inventive step** requirement involves an assessment of what would have been obvious to an unimaginative but skilled person in the art. Where the invention occurs at the level of the human designer or operator of the AI system, this can be analysed in the usual way. However, if the system develops so that the AI itself is deemed to be the inventor, this standard becomes difficult to apply. Determining the appropriate level of “skill” and “inventiveness” in order to make this assessment will not be a straightforward task in relation to an AI system. We discuss this further under Q13.

The **enabling disclosure** requirement involves disclosing how the invention works so that it can be carried out by a skilled but unimaginative person or team. Again, this assessment is difficult in the context of AI inventions.

First, disclosure of exactly how the AI system operates may not be possible because it is not fully understood by its developers and operators. Disclosure of information about the system and the input data used may not be enough to replicate its functionality.

There are analogies from the history of patenting. When chemical compounds were made before the era of modern chemistry, it was not always possible to provide a chemical formula for the compound, or describe the chemical reaction. In that case, the compound could be claimed as the result of using certain starting materials, subjecting them to certain process steps, and providing characteristics of the output. This concept is still recognised in the permissibility of product-by-process claims, where the product itself may not be possible to characterise completely.

By analogy, it may be possible to claim the result of an AI operation by describing the inputs, the externally identifiable characteristics of the AI operation, and features of the end result. An apparent difficulty with this is that the data used (the training data) will usually affect both the operation of the AI engine and the output. In theory, this might be resolved in a similar manner to inventions involving biological material, by having a database deposit. However, this itself is likely to create significant issues in terms of the volume of data, its security and confidentiality, and data privacy issues.

Second, the test for a sufficient or enabling disclosure includes a consideration of what the ordinarily skilled person or team needs to be able to perform the invention claimed.

A substantial amount of patent litigation takes place within the life sciences sector, with comparatively few disputes involving patents for computer-implemented inventions. Numbers of disputes are likely to increase as more patents proceed to grant, and the AI field becomes increasingly crowded. Litigation, although often undesirable for the participants, is helpful in that it shines a light on the legal requirements for patent validity.

In the life sciences field, many court rulings have addressed what amounts to an enabling disclosure. For example, the recent Supreme Court ruling in *Regeneron v Kymab* illustrates the extent of disclosure required to provide support for broad patent claims. These rulings demonstrate that a substantial burden rests on the patentee in providing sufficient information in the patent. The disclosure must be enough to enable skilled people reading the patent to carry out its teaching across the scope of the patent claims. We discuss this further under Q.11.

9. How difficult is it to secure patent protection for AI inventions because of the list of excluded categories in UK law? Where should the line be drawn here to best stimulate AI innovation?

The list of excluded categories do present difficulties in relation to AI-related inventions, in particular the exclusions of mathematical methods in s1(2)(a), and schemes, rules or methods for performing a mental act, playing a game or doing business, or programs for computers in s1(2)(c). Navigating around these can involve contortions in claim drafting in order to frame the invention in a way that does not fall foul of the exclusions.

One option would be to remove the problematic restrictions altogether, falling back on the requirements for inventiveness and industrial applicability as adequate to define what should be patentable. This would leave it to the Intellectual Property Office and the courts to determine how the rules apply in each situation and to flesh out general principles from individual cases. Although this method would be able to evolve and adapt, it does present potential uncertainty to both patent applications and potential infringers until a body of case law is established. It also requires parties to engage in office proceedings or litigation at their own expense.

Another option would be to set out in guidance what categories of innovation are considered to fall within the excluded categories. This would need to evolve alongside developments in the technology.

There are several different rationales for excluded matter categories. However, a core principle is to exclude areas where claiming a monopoly is most likely to have an adverse effect on innovation, or to encroach on human behaviour that would be inappropriate to subject to the controls of patent protection. Examples are methods of performing a mental act, and methods of doing business. Experience in the USA, which for many years permitted some business method patents, has shown that patents covering minor innovations were used to extort substantial payments with no indication that they had contributed to innovation in the relevant field.

Our contacts express concern that a similar issue could rise in relation to the use of AI. The scope of excluded and non-excluded subject matter needs to be considered carefully to ensure that it is likely to promote innovation rather than produce anticompetitive results.

10. Do restrictions on the availability of patent rights cause problems for ethical oversight of AI inventions?

We do not consider publication for the purposes of ethical oversight to be a role for the patent system to fulfil.

The disclosure obligation in the patent system is there to enable others in the field to understand and build on the patented invention. Ethical oversight of AI inventions is better addressed in other legal contexts. For example, obligations to comply with data privacy law would require developers to build in transparency in situations where data about individuals is processed.

11. Does the requirement for a patent to provide enough detail to allow a skilled person to perform an invention pose problems for AI inventions?

In answering this question, you may wish to consider:

- *is it clear how much information would be sufficient for a skilled person to be able to work the invention?*
- *could there be uncertainty knowing when an AI could be obtained by a skilled person to achieve the specific purpose of a patent claim and when an AI would need to be specified in a patent application?*

- *what are the consequences if the details of AI algorithms need to be disclosed?*
- *if AI is making decisions in a black box:*
- *Could there be a need to disclose more than a basic trained AI model, for example training data or the coefficient or weight of the model? If yes, is it clear how much information would be sufficient for a skilled person to be able to work the invention? Are special provisions needed for this information to be filed and stored?*
- *What would be the effect if competitors could use data to quickly train a different AI model?*
- *How would the skilled person know whether the invention could be repeated across the breadth of the patent claims or whether a claimed result could be achieved?*

As discussed above, the requirement to provide sufficient disclosure to enable the invention to be performed across the claim scope presents a challenge for AI inventions. Because a body of case law has not yet been established in this field, there is little judicial guidance on the amount of disclosure that will satisfy the enablement test. However, case law from other fields, such as biotechnology, can be applied by analogy.

The Supreme Court's recent decision in [Regeneron v Kymab](#) sets out the test for an enabling disclosure in a biotechnology context. In that case, the Court of Appeal [decided](#) that a principle of general application had been disclosed in the patent specification. As a result, broad claims covering genetically engineered mice able to produce humanised antibodies to almost any antigen were allowed. The Supreme Court, however, took a different view. Because of the additional work that would be required by a skilled team seeking to implement the invention beyond a very limited range, Regeneron's broad claims were disallowed.

We envisage similar difficulties in applying the sufficiency test in an AI context. If it is unclear exactly how the AI is operating, then providing enough disclosure to enable others to work the invention may not be possible. In addition, we think that it will be difficult to provide disclosure that is extensive enough to support a useful claim scope. One or two examples of how the invention is put into practice may not go far enough to support broad claims that protect the wider applicability of the innovation. However, until more patents are tested through litigation it will be difficult to assess this accurately.

For inventors, especially at the early stages, this presents a difficult choice. Should they apply for patents with broad claims, and include extensive disclosure so that a validity attack is less likely to succeed? Should they make a limited disclosure with the risk that any claims granted may be narrow or subject to attack? Or should they avoid patenting altogether because of the risk of extensive disclosure for little or no patent coverage?

In reality, the existence of plausible patent applications is likely to offer a degree of competitive advantage, whether or not those applications eventually result in broad claims, or whether or not those claims would survive a challenge to validity. By the time opposition proceedings or litigation has run its course, the technology may well have moved on. But is this really the way that patent protection should be used? Or should a shorter term, more focused type of right be available?

Overall, we consider that a new kind of right does offer some advantages. However, there are also potentially serious disadvantages, such as the prospect of oppression of smaller players, and in increasing clutter of proprietary rights. In reality, is innovation driven more by skills in one organisation

moving to another, rather than publication in patents? New innovations can “catch fire” without the need for the formal teaching in documents such as patents.

12. In the future could there be reasons for the law to provide sufficient detail of an AI invention for societal reasons that go beyond the current purposes of patent law?

This is likely to be the case for many areas of AI innovation, such as software-based medical devices, and employee selection systems. However, we do not consider patent law to be the right place to address this.

Specific law and regulation to address the need for AI transparency in situations where this is thought to be necessary would be more appropriate.

If, for ethical or other reasons there is a requirement for disclosure in relation to AI inventions, adequate patent or other IP protection may become significantly more important. Many of our contacts currently rely on the protection afforded by confidentiality. This might no longer provide adequate protection. There is some analogy here in relation to data exclusivity for medicinal product authorisation data, although it is not clear that a directly analogous approach would work in relation to AI.

13. Does or will AI challenge the level of inventive step required to obtain a patent? If yes, can this challenge be accommodated by current patent law?

The inventive step element of patentability focuses on the human inventor. While AI systems are carrying out specific tasks designed by human developers, then this standard could be assessed with reference to the skilled team of AI developers or users. This can be accommodated by existing patent law.

As AI systems move towards artificial general intelligence, the inventive step test no longer makes sense. We can no longer turn to the skilled individual or team as the benchmark for skilful but non-inventive development. AI systems are likely to improve rapidly and present surprising advances in capability. This means that there would be no standard benchmark available, unless a detailed assessment of the notional average system at the time the patent was filed could be established. We do not consider this to be practical at present. Nor do we consider it to be appropriate, especially in the light of our comments as to the AI as inventor.

Instead, it would be necessary to introduce either a new test for the inventiveness requirement of patentability, or establish a new form of protection for AI developments.

14. Should we extend the concept of “the person skilled in the art” to “the machine trained in the art”?

In relation to inventions made where AI is a tool – i.e. artificial narrow intelligence, the appropriate test is that of a skilled team, and where such a team would have recourse to the use of AI.

Where an invention is made by an AI agent (which, based on the discussion above, we see as being when (and if) artificial general intelligence is developed), the concept of “the machine trained in the art” is superficially attractive. However, we do not regard this as a sufficient change to reflect the fundamental changes that will be necessary.

The “machine trained in the art” concept envisages an AI system that operates within a defined sphere addressing a defined problem, such as dosage regimes for pharmaceutical products. The “machine trained in the art” concept would, in our view, be far too static and constrained to be useful in a fast-evolving field. We expect that new approaches will be developed that will find new ways of addressing

particular problems, taking in new data sets. The notional “trained machine” will not be useful as an inventiveness test.

While AI systems operate on defined problems using defined data, the relevant test can still usefully look to the individual or team of individuals that develop or use the system. Once general AI without methods and data defined by humans is used, a new system will be required. At that stage, it would seem more appropriate to set a new test based, perhaps, on a concept of “material improvement” over what has previously been made public.

15. Who is liable when AI infringes a patent, particularly when this action could not have been predicted by a human?

If the infringement clearly depends on decisions taken by the developers of the AI, and is foreseeable from the functioning of the AI system, then we suggest that the developers should be considered responsible for the infringement.

Where the actions of the AI lead to patent infringements in ways that could not have been predicted the application of infringement rules are more difficult.

Normally the operator of a product or process that infringes would be liable. In addition, a supplier of that product or process would be liable, either directly or indirectly as supplying an essential means. It is also possible that the infringement could arise because of the structure of a dataset used to train an AI. For these aspects, the ordinary rules of infringement do not need to be adjusted.

However, if the product or process involves an essentially hidden and/or transient process, it may not be possible to determine whether or not it infringes. This is not a problem unique to AI products or processes. Any process patent (or product by process patent) may run into a similar difficulty, although the problem may be greater for AI given its complexity. The skill of the patent attorney would then be needed to draft a patent so that it includes readily provable elements

A further difficulty may arise in relation to innocent infringement and infringement of process claims:

- a) no financial remedy is awarded where an infringer did not know or had no reasonable grounds to know that the patent exists. This, of course, gives rise to the desirability of providing notice of patents on a product or documentation, and on the proprietor’s website. However, if the user does not know (and has no particular interest in) the specific design of the AI system, or the steps it uses, that user may have no knowledge of any patent. Any knowledge they have may be completely meaningless, and a court might accept that an essentially hidden process might be caught by this exception; and
- b) infringement of a process claim requires that the potential infringer knew or should have known not only that the patent existed but also that the process infringed. This is difficult for an infringer to prove, and even more so if the user does not know what processes the AI is using, or the AI adapts its processes/algorithms depending on its input data.

One option would be to deal with liability for infringement by an AI through insurance. A compulsory insurance scheme would impose a “tax” on innovators and would be a disincentive for innovating. However a voluntary insurance scheme might help to address the problems of unintended infringement.

A second, and we think better, option would be a system of royalty payments where infringement takes place in a way that could not have been foreseen by the system’s developers.

However, there is a danger in developing new ways of defining liability for in respect of breaches of intellectual property rights carried out by AI systems as a stand-alone exercise. There are, of course, wider questions about liability arising from the “wrongful” actions of an AI system (torts such as libel, breaches of contract etc). To address these independently from each other runs the risk of inconsistency and conflicting requirements for developers and operators.

16. Could there be problems proving patent infringement by AI? If yes, can you estimate the size and the impacts of the problem?

The issues identified in the Consultation around understanding whether a “black box” process is infringing and where the process has taken place present real difficulties, especially as the technology develops.

If implementers understand how their AI system operates, the use of a “product description” or “process description” in the usual course of litigation can be used to help the court to assess whether (and indeed where) an infringement has taken place. However, once the technology has developed to the point where this can no longer be provided proving patent infringement will become very difficult.

This may, in part, be overcome by a patent attorney drafting patent claims to cover readily provable elements. However, such a constraint may mean that the applicant does not claim the invention in a way which best represents their real contribution.

Copyright and related rights

1. Do you agree with the above description of how AI may use copyright works and databases, when infringement takes place and which exceptions apply? Are there other technical and legal aspects that need to be considered?

The summary in the Consultation is, of course, at a high level and does not give a complete explanation of copyright law. It gives a useful overview of the issues.

We do not consider comparing the AI system to the human brain to be helpful. The AI system should be treated like any other computer technology. Bringing in this analogy with the human brain, we believe, adds confusion.

With AI innovations the data itself is of key importance. Many organisations approach their data assets with great care, and impose restrictions on access. Healthcare data, such as that generated by the NHS, has great value. Controls on access to and use of this data is an ongoing concern, and holders of the data are looking to achieve a greater return on use by others. Academic and public sector use may be permitted without charge, but where commercial exploitation creates value for others, the holders of the data have a legitimate interest in seeing a share of this value. The national character of the NHS does offer potentially very important advantages as compared to more fragmented healthcare systems in other countries. Although coordination of these data assets, and addressing issues around anonymisation of personal data, so that they can be used seamlessly, remains a challenge, there is an important potential value stream here that should benefit from appropriate protection.

In the connected transport field, data gathered from vehicles and road infrastructure has value, but also needs to be shared for the purposes of safety, efficiency, accident investigation, etc. The use of novel mechanisms such as data trusts can provide ways of sharing data while recognising proprietary rights.

As we have mentioned above, reliance on confidentiality is currently widely used to protect data assets.

We note that protection for databases is not straightforward. Although copyright protection may be available for databases that can be regarded as the author's own intellectual creation, it is not available for all databases. The sui generis database right is available to protect the investment aspect of database creation. However, it has serious limitations and is not widely relied on. It is an EU creation, and is not available in, for example, the US. Given the importance of training data to the development of AI systems, we think that this review offers an opportunity to review and clarify the law on protecting databases.

At present, the limitations on the protection of databases, and differences in protection between different jurisdictions, discourage sharing. Instead, holders of data may choose not to make it available to third parties or to make it available only on strict terms controlling use. Where a collection of data does not amount to a legally protectable database, contractual restrictions may, ironically, be tighter than would be permitted if the data were classed as a database. (Some of the limits on contractual terms applicable to databases do not apply if the data set is not a database.)

There are clear advantages to users of AI in having access to larger datasets. This needs to be considered in a review of the protection in this area. Database protection is not harmonised outside the EU. As a result, there are opportunities for the UK to innovate, both in terms of protecting datasets and mandating access. However, care must be taken to ensure that does not lead to holders of datasets isolating the UK from access.

Where copyright and/or database protection is available, there will be different views on whether it is necessary to provide for greater flexibility in access to and use of AI training data by third parties. Those who have invested in building databases will not wish to see others being able to access these without control or reward. Conversely, others will regard access to training data, particularly where it is held by a dominant player, as a necessary part of fair competition.

Finding a fair balance between these two positions will be challenging but, in our view, is important to resolve. A system permitting access to data on standard terms, like the Fair, Reasonable and Non-Discriminatory (FRAND) terms used in Standard Essential Patent Licensing, may be appropriate.

2. Is there a need for greater clarity about who is liable when an AI infringes copyright?

As with patent infringements (discussed above), the identity of an infringer is more straightforward when an AI system is carrying out defined operations on a defined set of data. Where AI systems become able to operate in ways that have not been defined by their developers, the picture becomes more complex. We consider that there needs to be further development of the law to determine whether copyright infringement in the normal way is appropriate, or whether some form of financial compensation (such as FRAND licensing) is appropriate.

However, we do not believe that the intellectual property system should operate in a vacuum. It is important that liability for infringement should be coherent with liability in other areas. For instance, liability for libel by an AI system is something that the law will need to address. These considerations are better addressed together, rather than different legal areas arriving at different and, potentially, contrasting approaches.

3. *Is there a need to clarify existing exceptions, to create new ones, or to promote licensing, in order to support the use of copyright works by AI systems? Please provide any evidence to justify this.*

It may be appropriate to introduce new exceptions dealing with the use of training data, particularly where the training data is controlled by a dominant entity.

As discussed above, licensing on FRAND terms may be an appropriate mechanism.

4. *Is there a need to provide additional protection for copyright or database owners whose works are used by AI systems? Please provide any evidence to justify this.*

As discussed above, existing copyright and database law already provides some protection for copyright works and databases used by AI systems. However, this is not a coherent form of protection, with different requirements for these two kinds of rights. Indeed, a substantial proportion of databases will benefit from neither.

This review offers an opportunity to review and improve protection for databases. The protection offered by copyright is difficult to establish and of very long duration, while database right although sometimes easier to establish, is shorter in duration and not available to non-EEA creators. As the importance of databases grows, the need for a clear and coherent form of protection is more apparent.

In addition, the use of AI facilitates use of copyright works to generate derived works that would not normally be regarded as reproductions of the whole or a substantial part of the original work, but still benefit from the creative effort applied by author. The question of what amounts to making a copy of the whole or a substantial part of a work needs clarification, especially where AI systems are engaged in making a derived work.

Consider, for example, a decorative design that has exhibited considerable commercial appeal. Another business uses an AI to derive a design. The new design includes those features of the original design which the AI has identified as being most appealing to buyers, in combination with additional elements which AI identifies as creating appeal. Clearly, the derivative work has benefited from the creativity put into the earlier work, but may not be classed as an infringement under current law.

Likewise, an AI might be used to analyse novels by a particular author and produce a work in a similar style. While the derivative work might not copy earlier works in a way that would amount to infringement under existing law, it clearly owes much to the creative work of the author.

5. *Should content generated by AI be eligible for protection by copyright or related rights?*

Where human creators use an AI system to develop works such as pieces of music, it is appropriate for those individuals to acquire copyright in the works produced, in the way discussed in the Consultation.

The concept of “computer-generated works” with no human author does not sit well with the rest of copyright law. This aspect of UK law may now need to be revisited. A new approach focusing on protection of AI-generated works that will reward the investment in developing and operating the AI system should now be considered. Bringing this into line with the “entrepreneurial works” such as sound recordings and films, instead of with works involving human creativity, would be appropriate.

6. *If so, what form should this protection take, who should benefit from it, and how long should it last?*

As discussed above, a more limited form of protection for computer-generated works similar to that provided for films, etc would be appropriate. Where possible, the UK should seek to align with other jurisdictions to offer a coherent landscape of rights. As noted in the Consultation, other jurisdictions

have not adopted the approach of affording specific copyright protection for computer-generated works and this is not seen as a particular advantage for UK innovation.

7. *Do other issues need to be considered in relation to content produced by AI systems?*

The broad and long-lasting protection provided by copyright should not be afforded to content generated essentially by AI systems.

8. *Does copyright provide adequate protection for software which implements AI?*

Copyright protection for software is widely relied on and regarded as adequate.

It may be appropriate to consider whether copyright protection is sufficient to protect languages developed for AI applications and data structures, although this has not been raised as a current need by our contacts.

9. *Does copyright or copyright licensing create any unreasonable obstacles to the use of AI software?*

As the Consultation points out, many basic AI tools are made available under open licensing arrangements. This provides an important stimulus to innovation.

It is, of course, open to developers to create their own software for internal use or to license on commercial terms. These mechanisms do not present unreasonable obstacles to the use of AI software.

Designs

1. *Do you agree with the analysis above which concludes that it is not possible for AI to be the author or owner of a UK or Community design?*

Yes.

2. *Are there, or could there be, any tensions with the current legislation when seeking to register a design or be recognised as the owner of an AI-created design? Who would be the legal entity applying for the rights?*

When a design is created by an artificial narrow intelligence system, the human developer(s) or operator(s) of the system should be recognised as the owner. If these are different entities, it may be necessary to recognise them as joint owners, although this can cause difficulties in practice. The law could define which person should be deemed the owner. Any agreement between them as to which is to own any designs produced should be recognised.

If and when designs are created by an artificial general intelligence it may become more difficult to identify an appropriate owner. In this situation, a system aimed at incentivising the human or legal entity that is responsible for the AI would be appropriate.

The RDA and CDPA set out that the person who made preparations for the creation of a computer-generated design is the author or creator (s2(4) RDA and s214 CDPA). If read broadly this provision could give the creator of the AI system rights to be identified as the author of a design, even if they have not inputted data required for it to operate.

3. *Who should be recognised as the author of a design created by AI where the system has been bought from a supplier, and the buyer has provided input or data to the system? Does the wording of legislation need to be changed?*

The wording of s.2(4) of the RDA and s.214 of the CDPA differs. RDA s.2(4) refers to the situation in which there is no human author, while CDPA s.214(2) refers to “a computer-generated design”. There is scope for these to be interpreted differently and they should be amended to be consistent.

As discussed above, a narrow artificial intelligence should not be recognised as creator or designer. Clarification around which person or persons should be recognised would be helpful. The AI developer and operator could be joint owners, for example, although this can present difficulties in practice. A better approach would be to identify as owner the person whose activity is closest to the production of the design – most likely the operator.

4. *Do you consider that legislation should be changed to allow AI systems to be recognised as the author of a registered design or designer of an unregistered design?*

No.

5. *If so, how should we assess when AI stops being a tool programmed by a human and becomes an intelligent entity capable of producing its own IP? What proof or evidence would be required?*

This is not an appropriate course. Design law is intended to protect and incentivise human creativity, and there is no need to incentivise the work of an AI in this way.

As aspects of the future UK designs framework will be based on concepts currently found in the Community Design Regulation, we would like your views on the following:

6. *Unlike UK domestic legislation, the CDR has no provisions relating specifically to computer-generated designs. Does this result in legal uncertainty in relation to authorship and ownership of computer-generated designs? Would the same apply to AI-generated designs?*

There is already legal uncertainty as between the provisions of the RDA and CDPA, and how they are to be applied. Clarification in this area would be helpful. The UK need not necessarily tie itself closely to EU legislation, which is likely to evolve in any event. However, a clear and consistent legislative approach is needed within UK law.

7. *Are there any other issues in relation to the CDR which we should consider in relation to AI?*

No comment.

8. *Can the actions of AI infringe a registered or unregistered design? Can AI do the acts set out in law (s7(2) RDA)?*

As with patent infringements, the identity of an infringer is more straightforward when an AI system is carrying out defined operations on a defined set of data. In this situation, it should be possible to assess whether the developer or operator of the AI has infringed rights to a design. Additional clarity on which person or persons should be liable would be helpful, focusing on the persons who are most closely involved with determining the output of the AI.

Where AI systems become able to operate in ways that have not been defined by their developers, the picture becomes more complex. Further development of the law would be needed to determine whether design infringement in the normal way is appropriate, or whether some form of financial compensation (such as FRAND licensing) is appropriate.

9. *When considering infringement are there, or could there be, any difficulties applying existing legal concepts in the registered designs framework to AI technology? Does AI affect the use of the “informed user” in measuring overall impression?*

The “informed user” test necessarily involves human input. This test could be applied to the work product of an AI to assess infringement.

10. *If AI can infringe a registered design, who should be liable for the infringement? Should it be the owner, the programmer, the coder, the trainer, the operator, the provider of training data, or some other party?*

Liability for infringement in relation to a design produced by a narrow artificial intelligence system should attach to the individuals or organisations most closely involved in the infringing activities. This is likely to be the operator of the system, and others dealing in the resulting work product.

If designs are created by an artificial general intelligence it becomes more difficult to assign liability. New structures may be needed to provide recompense to the owners of infringed designs, such as insurance on the part of the AI developer or operator, or a reasonable royalty approach.

Trade secret law

1. *Is trade secret protection important for the AI sector? Does the nature of AI technologies and business influence your answer?*

As highlighted in the Consultation, there are a number of difficulties for AI developers in obtaining and enforcing reliable IP protection. Many organisations choose to maintain secrecy as an alternative approach.

Patent protection requires detailed disclosure likely to include details of both training data and AI algorithms. Protection of inventions based on AI technology through the patent system is uncertain and may be difficult to achieve. Protection for databases through copyright and database right is not comprehensive, and there are several permitted exceptions to infringement. Copyright protection is available for software, but concepts and methods are not protected and competitors may be able to “design around” protection.

Currently, therefore, trade secret protection is important for many AI developers as an alternative method for protecting both databases and algorithms. It may be useful at defined stages of a project, for example while patent applications are being fleshed out, or to protect particular assets, such as key proprietary databases or methods or working that are difficult to protect using patents. Indeed, our discussions with business contacts indicate that secrecy or trade secret protection is the primary form of protection employed at present.

One school of thought sees an opportunity missed in over-relying on trade secret protection. By promoting a “first-mover” technology widely, so that it becomes universally adopted, greater benefit may be gained than maintaining its secrecy.

2. *Does the nature of AI pose any problems if UK trade secret protection is required? Does UK trade secret law give adequate protection to aspects of AI technology where no other intellectual property rights are available?*

Although trade secret protection has clear weaknesses – most obviously, the need to maintain confidentiality and the difficulty of retrieving the situation once material has been disclosed – it does provide a well-characterised form of protection. Indeed, it may be chosen above patent filings, or reliance on copyright or database protection, for the reasons discussed above.

Trade secret protection may not be suitable, however, where disclosure is obligatory. For example, disclosure can be required by data protection law, or for regulatory reasons such as product safety accreditation. Some form of auxiliary protection, akin to pharmaceutical data exclusivity may need to be considered if AI innovations are to achieve their full potential in regulated fields.

3. *What are the advantages and disadvantages of using trade secrets in the AI sector? Could information that is not shared inhibit AI development?*

The advantages and disadvantages of using trade secret protection will depend on your perspective.

For an organisation relying on this form of protection, advantages will be greater control of key assets and for an indefinite period. Disadvantages will include the need to maintain rigorous secrecy precautions with the associated risk of mistaken or malicious disclosure, and the lack of any monopoly protection if others can replicate your approach themselves.

For a competing organisation, reliance on trade secret protection will mean that access to databases or methods in ways that are permitted by IP law are not available. So, for example, text or data mining for non-commercial research, and experimentation with an invention will not be possible.



Trade secrets will remain attractive to developers where the IP system fails to provide appropriate protection. Adjusting the IP system to provide more comprehensive and predictable forms of protection for databases and algorithms is likely to reduce reliance on this form of protection.

4. *Do trade secrets cause problems for the ethical oversight of AI inventions?*

Although trade secret protection means that innovations are kept out of the public domain, we do not consider this to be a good reason to alter this form of protection. If developers take the view that the protection currently provided through formal IP rights is inadequate then they should be permitted to make use of trade secret law as an alternative. It would be unfair to single out this area of technology as unable to benefit from the use of trade secret laws.

Questions around the ethics of maintaining secrecy should be addressed elsewhere, where a need to improve transparency of AI systems is identified. For example, where AI systems are used to process personal data – facial recognition, or health information for example – data protection law and guidance is a more appropriate tool to address transparency.

Should you require more information on these responses please contact [REDACTED] at Mills & Reeve LLP.



Mills & Reeve LLP is a limited liability partnership authorised and regulated by the Solicitors Regulation Authority and registered in England and Wales with registered number OC326165 and VAT number GB 104 8345 88. Its registered office is at 24 King William Street, London, EC4R 9AT, which is the London office of Mills & Reeve LLP. A list of members may be inspected at any of the LLP's offices. The term "partner" is used to refer to a member of Mills & Reeve LLP.

The contents of this document are copyright © Mills & Reeve LLP. All rights reserved. This document contains general advice and comments only and therefore specific legal advice should be taken before reliance is placed upon it in any particular circumstances. Where hyperlinks are provided to third party websites, Mills & Reeve LLP is not responsible for the content of such sites.

Mills & Reeve LLP will process your personal data fairly and lawfully in accordance with professional standards and the Data Protection Act 2018, General Data Protection Regulation (EU) 2016/679 (as applicable) and any other applicable laws relating to the protection of personal data and the privacy of individuals. You can set your marketing preferences or unsubscribe at any time from Mills & Reeve LLP marketing communications at www.preferences.mills-reeve.com or by emailing preferences@mills-reeve.com

T +44(0)344 880 2666
www.mills-reeve.com
