Artificial intelligence and intellectual property

Response to call for views

In response to the UK IPO consultation, we set out below some comments from members of the UK IP team at Herbert Smith Freehills LLP in response to the call for views. We have been increasingly thinking and writing about the IP implications of AI in recent years and the consultation is a welcome opportunity to consider these issues and how they impact IP policy.

The views and opinions expressed here are personal views of the members of the UK IP team at Herbert Smith Freehills who have contributed to this response, and are not the views of Herbert Smith Freehills as wider firm, and nor are they reflective of the views of any of its clients.

Authors:

Andrew Moir

Peter Dalton

Priyanka Madan

Laura Adde

Elina Kysulchuk.

We have focused on the copyright and patent questions for the purposes of our response.

I. COPYRIGHT

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?**

   We broadly agree with the description of how AI might use copyright works.

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

   Yes. The current law is likely to lead to unclear results in the future.

   Under the Copyright, Designs and Patents Act 1988 ("CDPA"), copyright is infringed when a "person" does an act restricted by copyright in respect of the whole or a substantial part of the work, either directly or indirectly, without the licence of the copyright owner. This would include doing an act via software, since software is simply the tool by which the infringement occurs and requires human direction to perform the relevant act.

   For the purposes of the CDPA, a "person" would not include AIs (which are in any event incapable of owning property so as to pay out on any damages claim). So, under the current law, it would be necessary to identify a person who is liable for the acts of an AI which infringes the copyright of another.

   There are two main ways in which AI might infringe copyright – through the inputs into the AI (i.e. AI training) or as a result of the AI's output – works created by the AI which could infringe third party copyright.

   AI training is currently mostly human-directed – i.e. the human selects the data which is fed into the AI. This may change in the future as AI systems increasingly have the capacity to self-learn, or where the machine-learning engine is simply pointed at "the Internet" and has the capacity to determine for itself what it wants to look at. For now, it is arguable that the human directing the training would be liable for any infringement in the process.

   Equally, in respect of outputs, up until now most "AI-generated" works have been subject to some human input in order to produce the finished result (for example, by collating AI-generated elements into a cohesive whole). As such there has been some human involvement sufficient that were the work created to infringe copyright, one might say that the human has performed the act using the AI as a tool. This is unlikely to be the case in the future. AIs are likely to reach a level of sophistication at which they may do the act in question without instruction or direction from a human.

This reflects the dichotomy inherent in the discussion of AI and IP; AIs may reach the level of sophistication to be capable of infringing IP, potentially independently and without any human involvement, but are incapable of having personal autonomy at law or indeed of owning property, and may never reach that level of sophistication. As such, they can do infringing acts of without being legally liable for them.

As such, in the future, it may be possible for an AI to infringe copyright, but for there to be no person on whom liability can fix because no human has "done" the act in question. In practical terms, if the infringing work was exploited by a person, this would be a new infringement and at that point, the exploitation would represent a further and new infringement, which could probably be attributed to a person – therefore the practical implications may be less severe than thought.

The issue is arguably more problematic when considering AI training however (input). If the AI is itself determining what materials to "ingest" in order to train itself, it may be infringing IP in those materials, but there would be arguably no human responsible – and further there would be no additional act of exploitation since there would be no "infringing" work output at the end – the AI would simply be more suited to its given task. This may make it difficult or impossible to fix any person with liability despite the owner or operator of the AI having profited from the copyright works through the training of the AI.

It may therefore be necessary to cater for this, for example by creating a category of infringement whereby the person for whom the generation of the work by the AI is performed (i.e. the owner of the work) is liable for any infringement caused by that work, or, in the event such a person does not exist, the owner of the AI itself. This would mirror the provisions at s.9(3) CDPA regarding ownership of computer-generated works. This would then be subject to any potential exception regarding these works (see below).

There is also an overlap with vicarious liability. Even now, when an employee creates a work in their course of their employment and in so doing infringes the copyright of another, in practice the aggrieved party would commonly seek redress from the employer company. We would suggest that the position would need to be similar in respect of AIs. In other words, a company using AI in the course of its business via its employees would ultimately be liable to the extent that AI infringed the copyright of another.

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.**

Yes. Current copyright exceptions, such as those at s28A and s29A would not apply to this activity and therefore thought should be given as to whether a new exception or other mechanism is desirable in order to ensure that AI development is not held back in the UK.

Any exception which is restricted to "non-commercial use" (or that does not specifically carve out training activities – see below) is likely to inhibit the development of AI systems, as it would be difficult to say that training an AI, even if the immediate tasks is non-commercial, is not a commercial activity if the AI will have the benefit of that training when encountering future tasks done on a commercial basis.

Further, if one accepts the premise that AI is not simply a "tool" but has some analogies to human input, restricting any exception to non-commercial use only becomes less defensible on a legal-theory basis, as humans are not restricted to using the learnings they have obtained from copyright works for non-commercial purposes. Similarly, as a matter of employment and confidentiality law, employees are not restricted from using information in their head in subsequent employment save to the extent it can be classified as a trade secret.

One option would be to create an exception for AI training where the party for whom the AI training is being performed is a lawful user of the work and as long as the work is not incorporated into any output of the AI in a manner which would, if performed by a human, infringe the copyright in the work (in other words, replicating the situation for humans studying lawfully obtained copyright works). The "lawful user" requirement would be similar to that present in ss.50B and 50C CDPA – in other words, it would not require that the

licence to the work include a specific right to use it for AI training, only that the user was a licensee. Such an exception would be bolstered, if there was also a section voiding any terms of a licence which sought to restrict this right, as is present in ss.50B and 50C of the CDPA. Particular attention might be needed in some contexts, for example, web-scraping for training purposes, to ensure that there was clarity as to whether that was covered or not.

An alternative would be to create a public mandatory licensing model, akin to PRS for Music for musical works, or the PLR for library works. This would make it easier for AI developers to access works for AI training without having to lawfully acquire them from the copyright holders. The AI owner could instead pay a levy to the public licensing body, which would then make payments to the owner of the work. This may increase the ease with which AI is developed and therefore the speed of innovation in AI, but is likely to reduce the royalties paid to copyright owners and come with an organisational cost which may need to be met from public funds. It would also be challenging to track the copyright works being used if the AI engine was being trained from resources on the Internet (e.g. from news sites). This model would be limited to use in training AI and would not cover incorporating copyright works into outputs generated by AI.

However, given the bureaucratic costs of this, the exception solution is likely to be preferable.

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.**

The impacts of AI models are unclear. In general it is likely to be harder for copyright owners to enforce copyright and other IPR against technology companies developing AI, especially where they are used as training materials, due to the difficulty of proving infringement, as outlined above. This points to specific provisions being added to the CDPA to deal with AI infringement and fix liability, as suggested in the answer to Question 2 of this section.

There are also problems associated with disclosure – it can be difficult to actually identify infringements when AI and other software operates in a “black box”. However, it is unclear that this can be readily dealt with by the CDPA as it is a wider question of civil litigation policy. In any event, the same problem exists for the human mind, meaning that the real litmus test will be what the inputs and outputs are to the system, rather than what happens within it.

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

The approach to this question to date has been to find that copyright in computer-generated programs belong to persons who .

The CDPA provides protection for computer-generated works, being those generated “in circumstances such that there is no human author of the work”. Section 9(3) of the CDPA further provides that the “author” of a computer-generated literary, dramatic, musical or artistic work shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken. Under the legislation, the author will own and benefit from the copyright.

On the face of it, these provisions already provide copyright protection to AI-generated content; however, these provisions were made over 30 years ago and there is very limited case law dealing with them. These provisions envisage computers as tools used by humans, not as autonomous creators. If AIs are characterised as tools then these provisions are likely sufficient and offer a pragmatic solution, however, we query the validity of classifying all AI systems as tools. If an AI could be said to be an autonomous creator, then we would query whether it is appropriate to extend the scope of the existing provisions to cover such AI.

For example, in *Nova Productions Ltd v Mazooma Games Ltd* [2006], it was held that individual frames shown on a screen when playing a computer game were computer-generated artistic works. The author of these frames was the person who had devised the rules and logic used to create them. In this way we can see that the law has found ways in which to find the necessary originality and impute it to a human author. However, AI would stretch this decision because there is a far greater distance between the human and the end output. Arguably it would not be a suitable decision to follow if the question was who owns a work generated independently by AI.

We consider that there are two aspects to the question of whether works generated by AI should be protected by copyright.

1) At a theoretical/'philosophical' level, should AI-generated content be eligible for protection?

Copyright is intended to protect and reward creativity, but it is an open question as to what that concept means when applied to AI. Following *Infopaq*, UK and EU copyright law requires copyright works to be the author's own intellectual creation (previously skill labour and judgement). These tests presuppose a human author and do not translate easily to algorithms.

Research is ongoing as to whether works generated by AI, particularly in scenarios with zero human involvement, are truly original or creative, from both a mathematical and metaphysical perspective – but we are unlikely to have a definitive answer in the near future, if ever.

It is important to note that many IP practitioners are of the view that copyright protection demands some level of human involvement. For example, in 2019 the AIPPI published a resolution to the effect that AI-generated works should only be eligible for copyright protection if there is human intervention in the creation of the work, and the British Copyright Council have stated that granting copyright protection to machines devalues the fundamental reason for copyright: protecting human endeavour.

However, there is an assumption here that the use of AI remains very much the exception than the norm. If one looks at the longer-term view, as the use of AI becomes more and more mainstream, it could come to pass that a large cross-section of business and commerce is deprived protection for its work.

2) at a practical level, do we nonetheless need to protect AI-generated content?

Pending answers to more fundamental questions about creativity in AI, in the meantime, and from a practical perspective, there is a need to foster innovation and investment in AI. Denying protection to AI-generated works may hinder that goal, so in our view, some form of protection is necessary, but that does not necessarily need to be copyright.

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

There is a suggestion that AI-generated works may merit an alternative form of sui generis protection, a possibility raised by both the AIPPI and the British Copyright Council. We would recommend further research into whether there is any need to create alternative rights for AI-generated works, or whether for present purposes it would be sufficient to declare that AI-generated works fall within copyright protection (whether expressly under the existing computer-generated works provisions in the CDPA by classifying AI systems as "tools", or under new provisions). However, we acknowledge that creating sui generis rights rather than extending the scope of copyright to expressly include AI works would provide a practical solution to the problem, without detracting from the value of "human endeavours", and without requiring answers to the challenging metaphysical questions mentioned above.

Who should benefit?

If AI is classified as a tool, then section 9(3) CDPA, which vests ownership of copyright in the author (in other words, the user of the tool), provides a practical solution to this problem. If the AI cannot properly be called a tool, because no human-made any "necessary arrangements" for the creation of the work, then the situation is likely analogous with employment law, and ownership of copyright should vest in the developer or licensee of the AI (whether a natural or corporate person). We do not think that AI systems themselves should own or benefit from IP rights so there is no need to create new artificial legal personalities and, as noted above, AI is in any event incapable of owning assets and thus could never accrue licensing fees etc. in its own right.

How long should it last?

Currently, computer-generated works are protected under the CDPA for 50 years (s.12(7)). We do not consider this issue as fundamental as those above – if AI works are to be protected by the computer-generated works provisions, we see no reason to reduce or extend this time period in respect of AI-generated works; if new sui generis rights are created, we would recommend gathering further views on the duration of protection. We would, however, point out that, save for a few "milestone" examples of AI-generated artwork, the value of AI-generated works is currently relatively minimal in comparison to the value of the underlying AI (as opposed to works with human authors, wherein the value lies in the work itself, "irrespective of quality").

There is a question as to whether there should be a correlation between the value of an AI-generated work and the length of protection it should receive – should a set of one thousand AI-generated radio jingles receive the same length of protection from reproduction etc., as a full-length bestselling novel written by an AI? However, determining how such a relationship should be legally defined is likely to be extremely difficult so on balance we would not recommend going down such a route.

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

We have discussed these in our answers above.

II. PATENTS

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

As in other technical fields, the patent system incentivises innovation and investment in AI development. However, we wish to note expressly the difference between patenting AI technology and patenting inventions generated (or contributed to) by AI systems. Whilst there are difficulties with the former, the latter raises a number of issues which are altogether more fundamental, which we have sought to demonstrate below.

2. Can current AI systems devise inventions? Particularly:

a) to what extent is AI a tool for human inventors to use?

AI can certainly be a tool, and has been used as a tool in the context of drug discovery in the pharmaceutical sector, for example, for quite some time.

Advanced AI systems which create inventions (as opposed to AI systems which automate parts of the process of invention for the benefit of human inventors, such as by identifying patterns in a data set or identifying potential therapeutic compounds from a class) may be difficult to characterise as 'tools'. In the DABUS case, for example, we understand that DABUS generated the two inventions, and identified the novelty thereof, autonomously. One might characterise DABUS at a high level as an advanced tool used for the purpose of generating new consumer products, but it may be difficult to argue that systems which invent with minimal or no human direction or input are really 'tools' within the ordinary meaning of the word.

Classifying all AI systems as tools may be a simple way to resolve some of the AI inventorship issues, although it is unlikely to resolve all the problems, particularly as AI systems develop in capability and autonomy.

b) could the AI developer, the user of the AI, or the person who constructs the datasets on which AI is trained, claim inventorship?

As the law stands, any of these parties may be entitled to claim inventorship (or joint inventorship), provided that they meet the existing requirements for inventorship under the Patents Act 1977 ("PA 1977") (i.e. they must be the actual deviser of the invention under s.7(3), as interpreted by case law). In other words, that party must have contributed to the "inventive concept" (*University of Southampton's Applications* [2005] RPC 11).

The "user" is the most likely of these three parties to be an inventor within the meaning of the Patents Act. This is because being a "user" suggests that they are using the AI as a tool – perhaps to optimise an existing design that the human user has created, for example. As such, the user is likely to be most closely connected with the ultimate invention and most likely the "actual deviser" thereof.

If the AI is a tool, then the developer may be considered analogous to a toolmaker. Provided that the toolmaker contributed to the inventive concept of the ultimate invention, then they, too, may be an inventor, or joint inventor. Inventorship would be most unlikely where the developer has no connection with the use of the AI, as it would be difficult to demonstrate that they contributed to the inventive concept, in the same way that proprietors of off-the-shelf software have no rights in patentable inventions created by users of that software. If the AI is not a 'tool' and is inventing autonomously, the question remains whether the developer of the AI is sufficiently connected to the ultimate invention to have contributed to the inventive concept thereof. We note that naming the developer of the AI as inventor was an express possibility raised by the judge in the recent High Court decision (*Thaler v The Comptroller-General of Patents, Designs and Trade Marks* [2020] EWHC 2412 (Pat)).

Whether a trainer or a person who constructs the training (or active) datasets may be considered to be an inventor will depend on whether there is a meaningful correlation between the selected data and the ultimate inventive concept which makes that person an actual deviser. This will be more likely where that person is closely connected to the use of the AI, and less likely where they train the AI in a hypothetical vacuum with no knowledge or foresight as to its ultimate use which yields an invention. It also means that the nature of the data itself might play a part – for example, data which was obtained empirically and in which database or other rights might subsist might be considered differently to third party data which is being used.

c) are there situations when a human inventor cannot be identified?

Potentially, yes, where the AI has invented autonomously and no human can be said to be the actual deviser of the inventive concept. This is likely to become more of an issue in the future as AI becomes more sophisticated.

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

In our view, no. This is partly in view of other issues that this would cause around patentability (see later answers), and in view of more fundamental issues as to what it means to be an 'inventor'. These issues are similar to those outlined above in relation to copyright, including, for example, difficult theoretical and philosophical questions as to whether AI systems truly "invent". As noted above, we are unlikely to be able to answer these questions in the near future, and a practical solution is required in the meantime.

We do not consider that identifying AIs as inventors is a viable option. Whilst this may be the most "accurate" option, inventorship under the current law accords certain rights, and so far as we are aware, there is no desire in industry, academia, or otherwise, to grant AI

systems rights or any form of legal personhood. So long as inventorship carries rights, (for example, the right to grant of a patent under s 7(2)(a) of the PA 1977), we do not consider that AIs should be listed as inventors.

However, if a distinction were drawn between inventorship and ownership (see Q6 below), so that (accurately) listing an AI as an inventor did not grant any rights to that AI, then listing AIs as inventors would present a more feasible option, subject to the patentability issues described below.

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?

This will depend on whether the UK IPO permits other parties, such as developers or AI 'owners', to be listed as inventors (which was raised as a possibility in the DABUS case), and whether those people *want* to be listed as inventors (DABUS' owner, for example, thought this would be morally wrong). If it becomes common practice for owners to be listed as inventors (i.e. the "AI as a tool" model), then patent applications will still be filed for inventions created by AI. As such, we do not consider that this would discourage the patenting of future inventions, or innovation using AI, in the majority of cases.

If it is not possible or not accepted practice to list other parties as inventors of AI-generated inventions (whether because no human contributed to the inventive concept or as a matter of policy), the impact would be that AI-generated inventions would be patentable "but-for" the fact they are AI-generated. Such inventions would therefore be either unprotected or, more likely, protected by trade secrets or confidentiality regimes. Again, this is unlikely to discourage innovation using AI per se, particularly as currently the value lies in the AI system itself rather than the inventions they create. For example, we do not know whether DABUS' inventions, of a search and rescue beacon and a food container, would have met the requirements for patentability, but would, in any event, have been unlikely to be high-value inventions.

However, if AI-generated inventions develop to the extent that those inventions become extremely important or valuable, then reliance on trade secrets rather than patents will result in a lack of disclosure and knowledge in respect of important technology, which is precisely the situation that the patent system is intended to prevent. Such a scenario would also be likely to have implications in respect of Standard Essential technology and FRAND licensing, and on competition.

Of course, the other practical possibility is that this would simply drive the use of AI underground. In other words, AI would still be used as a means of invention but just not declared on the face of the patent. Given that, in many cases, to be sufficient a patent only needs to explain how the person skilled in the art can practise the invention (as distinct from how the invention was devised), this might drive a practice of challenging entitlement to patents as a defence to infringement.

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

Arguably the case is better expressed the other way as AI systems have no rights which can be infringed and as such it is difficult to ascribe moral implications to recognitions given to AI per se. Rather, the question might be whether it is immoral to name a person as inventor when in fact they had no real involvement due to invention having been invented by an AI. In the DABUS case, Dr Thaler felt that it would be morally wrong for him to take credit for an invention that was not his. This would be true whether the real inventor was a human or not. It is also a question of whether it is *inaccurate*, as opposed to immoral, to list as an inventor someone who did not contribute to the inventive concept.

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

We consider this to be one of the fundamental problems with naming AI as inventors under the current law.

As explained in detail in the High Court DABUS decision, inventorship currently carries numerous rights, including the right to grant of a patent under s. 7(2)(a) of the PA 1977. Only the inventor themselves can transfer this right to another. So, as AI systems cannot hold rights, nor transfer them, the current law precludes *anyone* from owning a solely AI-generated patent.

We leave open the possibility that inventorship may be divorced from ownership (and other rights) in respect of AI-generated inventions. In such a scenario, we consider that new provisions could permit the listing of an AI as inventor, whilst vesting ownership of the patent and all other rights therein, in a natural or legal person.

As to who that person should be must turn on the facts, but a model similar to employment law could be used. For instance, where the developer or "owner" of the AI is an individual, they should logically own the patent, but where that person developed the AI in the course of their employment, then the right to the patent should sensibly vest in the employer, in the same way that any other IP that the employee generates in the course of his or her employment belongs to the employer. Where an AI is licensed out to third parties, ownership of any resulting patents could be provided for contractually to vest in the licensee.

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

Patenting the AI system itself

So long as the AI system satisfies the criteria of patentability, i.e. requirements of novelty, inventive step, sufficiency of disclosure and industrial applicability, and not falling within excluded subject matter as a mathematical method or computer program, the AI system should in itself be patentable. However, in a seminar delivered by the UK IPO in January 2020, the UK IPO explained that AI inventions would typically fall into three categories:

1. Applications of AI to particular fields or problems;
2. Techniques for training the AI; and
3. The core AI itself.

With respect to the third category, the view was that such inventions, absent any technical application, are more problematic to patent. Specifically, the core AI algorithm may be considered a mathematical method and so under UK practice would be examined by the UK IPO in line with the EPO Guidelines for Examination G-II 3.3. AIs may well struggle to overcome the technical effect / contribution test as their effects largely exist within software itself.

Patenting inventions devised by, or using, AI

An AI cannot at present be named or considered as an inventor of a patent in the UK. The UK IPO's Formalities Manual specifically states that an 'AI Inventor' is not acceptable as this does not identify 'a person' which is required by law. Further, if the patent application is not amended to include the name of a natural person as the inventor, the application is considered withdrawn under s. 13(2) PA 1977. The December 2019 DABUS patents decision of the UK IPO, as well as the subsequent High Court decision on the same patents also confirms this position.

With respect to inventions made using AI, an issue is to decide whether the contribution made by the human person using the AI is sufficient for them to be an actual deviser of the inventive concept. The inventor is the person that formulated or contributed to the formulation of the inventive concept (*University of Southampton's Applications* [2005] RPC 11). This assessment will therefore turn on the nature of the invention and facts surrounding the specific contribution in each case. Some potential scenarios are:

1. Using AI to perfect a product or a process, which could not be done without assistance from the AI;
2. Using AI with no particular purpose or goal in mind, but building inventions on the basis of what the AI finds (e.g. if an AI processes data and discovers useful relationships, which are then used by the human inventor to make an invention. Here, the invention may not have been possible without AI).
3. Using AI to generate multiple outputs, with human intervention at the end of the process to choose which output is the best.

In all of the above situations, it is important to establish the level of contribution by the AI, and whether the contribution made by the human person is sufficient for them to be considered an inventor. Legislative or judicial guidance on these issues would be welcome.

Beyond issues of inventorship, an invention made by or using AI should also be patentable in that it also satisfies the requirements of novelty, inventive step, sufficiency of disclosure and industrial applicability. Here, there may be problems with respect to determining inventive step and sufficiency, e.g. if the AI is inventing in a "black box" type set-up, or if disclosure of the workings of the AI is necessary to put the invention into effect. That said, all that is needed is that the invention itself can be practised from the disclosure; one may not need to disclose *how* the AI invented it.

An invention using some level of AI (or crated by AI) should also not fall within the categories of excluded subject matter. As discussed in further detail below in response to Q.9 below, the use of AI does not by itself mean that the invention falls within excluded subject matter, but here the level of AI contribution needs to be assessed.

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

We do not necessarily see new problems arising (or at least it is difficult to currently predict them) However, as with other areas of IPR policy, as AI technology develops it is likely to become more autonomous and self-sufficient, and therefore the problems identified above are likely to become starker and more widespread.

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?

Generally, if the invention is an AI system, or AI is required to implement the invention, that does not by itself preclude the patentability of such inventions. For inventions that require at least a minimum amount of AI implementation (or a minimum amount of computer implementation) to work the invention, it must be determined whether the AI contribution would be excluded from patentability pursuant to the exclusions contained in s. 1(2) of the PA 1977, which in particular excludes mathematical methods and programs for computers.

Numerous cases in the context of computer-implemented inventions have attempted to clarify the scope of s.1(2) PA 1977. In *Aerotel v. Telco; Macrossan's Application* [2007] RPC 7, the UK Court of Appeal laid out the following test to assess whether the claimed invention (which involved computer implementation) was patentable:

1. Properly construe the claim;
2. Identify the actual contribution;
3. Ask whether the identified contribution falls solely within the excluded subject matter; and
4. Check whether the actual or alleged contribution is actually technical in nature.

In *Symbian Limited v Comptroller General of Patents* (2008) EWCA Civ 1066, the Court of Appeal held that a computer programme could not be excluded from patentability if it had the knock-on effect of improving the technical functionality of the computer, as a matter of

practical reality. In *AT&T Knowledge Ventures LP and CVON Innovations* [2009] EWHC 343, the UK High Court set out "signposts" as to what might constitute an "irrelevant technical effect", i.e. a technical effect that would fall completely within excluded subject matter – these signposts would be relevant in assessing AI implemented inventions as well.

Please also see the response to Q.7 re the UK IPO's views on patenting the core AI.

Overall, the difficulty of securing a patent for an AI invention very much depends on the way in which the AI is used in the invention, but the use of AI is not an automatic bar to obtaining a patent on the basis of excluded subject matter, just in the same way that a computer program is not an automatic bar at present. It is also difficult to draw a fixed line at this point, and perhaps not necessary to do so, as the issues of AI and patentability are only just beginning to emerge.

We would note in passing that the case law around the patentability of computer-implemented inventions is fraught with complexity and inconsistency in any event. Where the line should be drawn for the patentability of computer software is a topic the Courts have grappled with for some time; the position will likely evolve to include AI but is unlikely to become clearer. Ultimately as a matter of public policy the granting of monopoly rights, with the incentives that these can bring in terms of investment, need to be balanced with the potential freezing effect of 20 year patent rights on innovation.

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

    [Not answered.]

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?*

    [Not answered.]

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?*

    [Not answered.]

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?**

    This depends largely on the AI in question, and the extent to which it is used in an invention. The inventive step requirement is found in Article 56 of the EPC (and is implemented in the UK in s. 3 PA 1977): "An invention shall be considered as involving an inventive step if, having regard to the state of the art, it is not obvious to a person skilled in the art." The Pozzoli test (*Pozzoli SPA v BDMO SA* [2007] EWCA Civ 588 sets out how this is approached.

    In our view, there is no need to change the core law in relation to the inventive step in relation to either AI inventions, or AI-generated inventions. In general, the law should not in our view create special categories applicable only to AI.

    In respect of the AI generated inventions, the "AI involvement in the invention" i.e. devising it, is in our view irrelevant to the patentability and how it should be assessed. It might affect the definition of the skilled person, but that is entirely objective (not subjective) as a test. The question is whether the skilled person, at the priority date of the patent, would have applied AI to the problem. If so, then that must be the standard and it is irrelevant in that context as to whether the inventors of the actual patent used AI or not. The converse is also true – the actual inventors could have used AI, but where such use was not

mainstream enough to be used by the skilled person as at the priority date, it should not be taken into account when assessing inventiveness after the event.

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

No. It would be premature to do so at this stage, and these concepts also bleed-across into our answers above: to extend the definition of "person skilled in the art" to a "machine trained in the art" implicitly accepts that machines have the capacity to invent. If that is the case, it then becomes inconsistent with AI being incapable of being the named inventor of a patent.

Extending the definition of skilled persons to include "machines" poses its own set of problems – patent law is grounded in the idea of the skilled person being able to explain his findings or course of action taken, which may not be possible if the courts and parties were having to deal with a "machine trained in the art" instead (and even if so-called "explainable AI" is successfully developed, it is unlikely that courts or others would be able to interpret the "explanations" without considerable technical expertise). There are also various other practical issues – what would be the "machine trained in the art"? Would it be another machine with a substantially similar algorithm and purpose? Who would be operating this machine, and if a human operator is required, what would be the point of the new definition?

It is important to note also that AI algorithms can be entirely generic. Their function is dictated by their learning/training (c.f. a human that is a concert pianist compared to a human that is a painter - they are both human, but is their training that provides them with their capability).

It's the same for AI, so it may instead be better to evolve the definition of "skilled person" to include "skilled person using an [appropriate] machine", which to a large extent is already done, for example, in highly technical cases involving biotechnology, where notional skilled persons are often assisted with appropriate machines in carrying out their analysis.

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

As a starting point, it seems logical that the responsibility for the AI should rest with the person who controls and enjoys the benefits of that AI – likely the developer or owner, or in some instances, the user of the AI. Further, ultimately most infringements occur when an infringing product is put on the market or a process is used in commerce. Such infringements would be attributable to a human infringer, regardless of whether AI was the means through which the infringing product / process was generated.

This is similar to the concept of vicarious liability in employment law, provided that the AI's infringing act was not akin to an employee acting on a "frolic of their own" (although the person in control of the AI would be able to circumscribe its action by coding specific limitations, so if they intended or were ordered to stop the device from infringing, they could do so).

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

Yes, there are foreseeable problems in this regard, both in terms of determining infringement on the basis of substantive law and practical issues. Two such issues that should be addressed by the lawmakers early on are:

1) The substantive issue of establishing infringement when part of the infringing act takes place outside the UK – with the increased use of servers in multiple jurisdictions and cloud computing, this is likely to be an important issue for patent holders and the developers of the allegedly infringing AI alike; and

2) The practical issue of establishing infringement when it is not possible (or very difficult) to determine what exactly the AI is doing.

In relation to 1), determining the location of the infringement, such that the jurisdictional requirement in s.60 PA 1977 is fulfilled is likely to be a concern, for example in cases where the AI performs part of the infringing process on overseas servers. The UK case law on the use of overseas servers has not been consistent so far. In *Menashe v William Hill,* the court held that the location of the servers being overseas was irrelevant, as the ultimate "user" was located in the UK. Similarly, in *Illumina v Premaitha,* the court looked at whether the process would be "substantially used in the UK". This jurisprudence perhaps indicates that the jurisdiction requirement in s.60 can be satisfied if there is a sufficient nexus in the UK (in the form of customers or natural persons at the end of the process being located in the UK), but it remains to be seen how this would be applied in the context of AI.

In relation to 2), it may not always be possible to even establish what precisely the AI is doing. This could be the case in situations where the AI performs in a "black-box" type set up. This problem can potentially be solved by tailored disclosure requirements in AI cases (for example, an AI Product / Process Description from the developers of the AI, that explains its functioning), but such disclosure is likely to be complex, lengthy, and highly technical. Any expert evidence upon the functioning of the AI is likely to pose similar problems. However, this only matters where the internal functioning is relevant to the patent. If it is only the AI's external influences on the environment that are the subject of the invention, then how precisely the AI works might be irrelevant.

The size and impact of these problems can be quite large, especially as the investment in AI increases and more AI inventions are brought to the market. Early legislative or judicial clarification on these issues is therefore of the utmost importance.