Haseltine Lake Kempner

Haseltine Lake Kempner LLP
Redcliff Quay
120 Redcliff Street
Bristol BS1 6HU
United Kingdom

t: +44 (0) 117 910 3200
f: +44 (0) 117 910 3201
w: www.hlk-ip.com

Intellectual Property Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

30 November 2020

For the attention of the Intellectual Property Office

I hereby provide comments on behalf of Haseltine Lake Kempner LLP in response to the call for views on Artificial Intelligence (AI) and intellectual property. Haseltine Lake Kempner LLP are a UK based firm of patent attorneys, trademark attorneys and IP solicitors, which also has offices in mainland Europe. The firm includes many experienced patent practitioners representing clients whose technology encompasses AI.

Patents Questions

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

Patents are traditionally seen as a reward for investment in technology development: in return for disclosure of the development, they effectively provide a monopoly for twenty years. Additionally, patents can showcase the development through patent publication and prove ownership, but the commercial benefit is primarily through patent grant. The patent system should be able to play the same role specifically for AI technologies, but this is not currently happening, at least in the UK, and also potentially at the EPO. Since many AI inventions are viewed by the IPO as falling within at least the exclusion of programs for computers as set out in The Patents Act 1977 Section 1(2)(a) and possibly also within the exclusion of mathematical methods as set out in The Patents Act 1977 Section 1(2) (c), they may never result in the patent rights that reward development. There is a similar picture at the EPO. While a limitation to a technical field may assist a patent applicant in obtaining a patent, if the AI invention can be limited to a specific implementation to avoid the exclusion, the resulting patent right can be so narrow as to be worthless.

That is, many AI applications at the IPO have a high chance of being published, putting the invention into the public domain, but never gaining the full benefit of grant. Thus currently, the patent system could play a role in encouraging the development and use of AI technologies, but does not play this role to the same extent as for other technologies.

We believe that if the IPO develops a strategy for protecting AI inventions with the merit that they deserve, this could have far reaching implications for UK business, and the UK as a whole. In other words, if the IPO were to allow AI inventions to be protected in the same manner in which other inventions have traditionally been protected, then this could attract AI businesses to the UK, and the research and development associated with that.

Registration No. OC344 156, Registered Address: Redcliff Quay, 120 Redcliff Street, Bristol, BS1 6HU, United Kingdom. Regulated by IPReg.

Qu. 2 – Qu. 6

We note the recent High Court decision in the matter of Mr Steven Thaler's application relating to the AI machine DABUS, and the extensive consideration in that decision of the question of AI as an inventor. We are not submitting additional comments on this issue.

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

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

These questions are considered together in the following discussion.

Current practice

Current practice for the examination of inventions involving AI may encompass, inter alia, issues of excluded subject matter, inventive step and inventorship. While each of these practice areas are considered in detail in later questions, the issue of excluded subject matter is particularly dominant when considering the grant of patents for AI inventions in the UK, and would appear to be the largest perceived hurdle to obtaining patent protection in the UK for AI inventions.

With reference to the Manual of Patent Practice (MoPP), Section 1.39.3 states: "Inventions involving artificial intelligence ("AI") are generally computer-implemented, so their assessment in respect of excluded matter will therefore mirror that of any computer-implemented invention". While the MoPP also notes that many inventions involving AI relate in some way to a mathematical method, the principal guidance for the consideration of AI inventions is that set out with respect to computer implemented inventions. This guidance rests on decided cases including Aerotel, Symbian, AT&T, and HTC, all regularly quoted during examination of AI inventions at the IPO. While a detailed treatment of the existing guidance is given below with reference to question 9 of the present consultation, the MoPP provides a useful summary as follows (emphasis added):

"If an AI invention is directed to a specific technical process outside of a computer or if it forms part of the internal workings of a computer, then it will likely not be excluded. However, where an AI invention is directed to an excluded process (e.g. a business method) and there is nothing more to it, it will likely be excluded. Likewise, if an AI invention is directed to a non-specific purpose it may also fail by way of encompassing excluded processes (see 1.15 & 1.19). AI inventions or algorithms that provide an improvement in programming will likely be excluded following Gale unless they make a technical contribution (see e.g. 1.29.4)."

Applicants seeking to protect an AI invention are thus offered two main routes to establishing that their invention does not fall under the category of excluded subject matter; the invention must either result in an improved computer (operating at the level of computer architecture), or result in a specific improved technical process outside the computer. While the route of limiting to a specific technical process outside of the computer is often favoured by applicants seeking to protect AI inventions, many AI inventions may be considered to fall naturally between these alternatives, resulting in a computer that is improved with respect to performance of a specific task, which task may be of use in any one of a wide range of technical processes. Such an invention is often referred to as an improved computer program, or improved AI, and judged under current guidance to be non-technical.

Technicality and AI inventions

Some of the defining characteristics of AI are its capacity for data analysis, the making of assumptions, learning, reasoning, and the making of predictions, often at a scale and depth that is

beyond the compass of alternative rules-based or other procedures. By their nature, the above characteristics are non-specific to any given technical process, and nor do they operate at the level of a computer architecture. However, the implication in the existing computer implemented invention guidance that AI inventions, when considered at a high level of abstraction and independent of use case, cannot offer a technical contribution would appear to fail to appreciate the immense technical advantages that can be offered to a wide range of processes by an innovative AI solution.

What might be considered “pure AI” inventions, with wide applicability to a range of technical processes, are particularly susceptible to rejection on the grounds of lacking a technical contribution. Pure AI inventions generally relate to new ways of performing AI: new core algorithms that are able to perform new tasks. Generative Adversarial Networks, in addition to many Reinforcement Learning algorithms and other AI techniques, might be considered to fall into this category. Pure AI inventions are in effect a tool, operable to control or improve a wide variety of technical processes. In the absence of a change caused at the architecture level of a computer, such inventions are restricted to being protected in the context of a specific use case or application. The considerable technical investment and inventive skill that is required to create the tool, and ensure that it is fit for purpose, is not sufficient to establish a technical contribution according to current law and practice.

To be usable, AI inventions must be deployable, that is they must do what is needed fast enough and reliably enough to be commercially and practically viable, while using available computation and memory resources. Such resources may be limited in many practical implementations, particularly in the context of the ever-increasing demands on communication networks, and the revolution in connected devices and distributed computing. Creating an AI invention that is both effective and deployable is highly challenging, the work of skilled experts.

While the existing guidance relating to computer implemented inventions makes clear that an improvement in programming cannot have a technical effect in and of itself, the commercial reality is that an improvement in programming can make the difference between an AI invention being highly valuable to a wide range of industries, or being completely unusable in any practical context. An improvement in programming can result in faster performance of a task, the use of fewer compute resources, increased reliability, a requirement for less training data, advantages in security or privacy of data, the ability to solve new types of problems etc. All of the above can have huge implications for a wide range of technical processes, but cannot confer technical character under existing practice.

Considering the key concept of Generative Adversarial Networks mentioned above, these tools are widely applicable to a range of realised and yet to be realised technical processes. Yet it seems that under current practice, they would be treated purely as computer implemented inventions, and following the MoPP and barring any restriction to a particular technical process, categorised as a computer program ‘as such’. A notional patent applicant for this key concept may thereby be denied a patent. However, while the network is undeniably computer implemented, the notion of whether a contribution of this nature is in fact a computer program is highly questionable. It could be instead advanced that it could better be characterised as a new way of computing altogether.

The fact that AI inventions that are non-specific to a particular use case do not have a clear path towards being recognised as technical in nature is discouraging to patent protection. If protection is limited to a specific process, such protection can be too limiting to be of commercial interest. Even when protected in the context of a specific technical process, experience has shown that if the inventive activity relates to how the AI itself is adapted to achieve the technical effect, the resulting patent application is highly likely to face significant challenges during prosecution relating to an alleged lack of a technical contribution. This seems to be the case even when corresponding applications at the EPO are not subject to such challenges.

Comparison between UK and EPO practice

The MoPP quotes Birss J in paragraph 9 of Lenovo (Singapore) PTE Ltd, “although in methodological terms the approach in the UK and the approach in the EPO may look different, in practice they reach the same result, at least usually.” While this may be the case for computer implemented inventions in general, anecdotal evidence suggests that for AI inventions specifically, the EPO is often a more successful forum for protection of AI inventions.

As an example of the disparity in result when seeking protection for AI inventions in the UK and at the EPO, we refer to GB2513105 (GB1), and its family member EP2973241 (EP1). GB1 and EP1 relate to an Artificial Neural Network for signal processing, and were recently examined at the IPO and EPO respectively, each application having identical claims at the start of examination. While the potential issue of double patenting was raised on the prosecution of GB1, examination was nonetheless carried out, allowing for a comparison between the approaches taken in the IPO and EPO, and the resulting outcomes. Following clarity and unity objections at the search stage of GB1, all claims of GB1 were rejected as relating to a computer program as such and/or a mathematical method. The exam report cites Aerotel, Symbian and the AT&T signposts, as updated in HTC. The applicant has not yet submitted comments on the examination report. In the EPO, the same claims were subjected to unity and clarity objections, and initially rejected as lacking either novelty or inventive step. Following several rounds of examination and oral proceedings the application has now been accepted. During the EPO examination procedure, rejections relating to a computer program, mathematical method, or lack of technical contribution do not appear to have been raised, either in the context of inventive step or as a reason for rejection of the application, and have not prevented patent protection from being obtained.

The patent family briefly discussed above is merely one example that we consider to be representative of a larger problem, namely a divergence in outcome when seeking to protect comparable, or identical, AI inventions at the IPO and EPO, even when such inventions are directed to application of AI in a specific technical process.

Conclusions with respect to Questions 7 and 8

In light of the above discussion, our view is that current practice for the consideration of AI inventions in the UK does pose difficulties for obtaining patent protection for such inventions. In addition, in the specific field of AI, there is evidence that consistency between outcomes under UK and European approaches to computer implemented inventions is not being achieved.

In this context we note that much of the essential UK case law for the consideration of computer implemented inventions was established many years ago, when AI was in its infancy. The case law has been developed in the context of computer implemented inventions in general, and its application to the consideration of AI inventions would appear to be hindering the development of patent protection for AI in the UK. Under current practice, AI inventions may be categorised as computer programs as such as a result of being computer implemented, rather than following direct consideration of the question: is this invention at its core a computer program?

It is of concern that if, as we suspect, established precedent in the UK, and the way in which such precedent is being applied to AI inventions examined at the IPO, is a disincentive to patent protection, the prospects for improvement of this situation are slim. In order to evolve the case law position regarding AI inventions, and particularly the blanket applicability of computer implemented invention case law to AI, an applicant will presumably be required to litigate and pursue an appeal to the UK supreme court. The financial investment and time commitment required to pursue such an appeal is a significant obstacle to evolution of the current law and practice for the treatment of AI inventions in the UK.

Qu. 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?

As discussed in the answers to questions 7 and 8 above, particularly for pure AI inventions, it can be very difficult to secure patent protection, or at least difficult to secure worthwhile patent protection, if the application must be limited to a specific use to enable grant. In other AI inventions, such as those adapting to a particular hardware configuration, as long as the IPO Examiner studies the invention carefully on the merits, rather than resorting to a standard citation of the Aerotel, Symbian and the AT&T signposts, as updated in HTC, the exclusions need not make it difficult to secure patent protection. So the answer is that the difficulty with the list of excluded categories depends on the individual invention as assessed under the standard practice at the IPO.

Looking at this in another way, it is the application of the case law that draws the line for the individual inventions, rather than the list of excluded categories itself. And as previously explained, these are court decisions which do not relate to AI, and it will be costly for the users of the system to try to change the precedents used in examination at the IPO and draw a new line.

In order to stimulate AI innovation, the line between exclusions and inventions should be drawn to recognise as technical, and therefore patentable, AI technology which is new and inventive but not limited to a specific use or hardware benefit. How can we achieve this outcome? Maybe it is worth returning to consideration of the exclusions themselves.

When the exclusions were first written into law, it was a time when computer programs were simple static administrative scripts, and the exclusion of a program for a computer justifiably applies to such scripts. In this sense AI models are neither simple, nor are they static or administrative. AI inventions can be more easily envisaged as a learning ecosystem of nodes. So they cannot simply be viewed as equating to a program for a computer.

Neither are AI models traditional mathematical methods: they are not defined by fixed equations but evolve in a way that is not predictable or pre-set to give an output from an input, with the relationship between the input and output being provided by internal learning which is unknown to the user who starts the model.

As a result of their very nature, AI inventions should therefore be viewed as falling outside the list of excluded categories and thus outside the case law which has developed around these exclusions. This realisation, if adopted into UK IPO practice, would allow AI inventions to be examined on their merits in the same way as for other modern technologies.

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

The patent system is not the correct vehicle for ensuring ethical oversight of AI inventions. Applying for a patent is voluntary, and ethical oversight (where necessary) should not be limited to those who choose to use the patent system. Thus the restrictions on the availability of patent rights do not cause problems for ethical oversight of AI inventions. The issue of oversight must be addressed elsewhere, in public law.

Qu. 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?

The standard for providing an enabling disclosure need be no different when considering AI inventions as opposed to other inventions. The precise amount of detail required may differ depending on the nature of the invention. The onus would be on the patentee to provide sufficient information and it would be open for a third party challenging a patent to argue that the standard had not been met. For example, established principles such as those outlined in Biogen that the extent of the monopoly claimed should not exceed that which is enabled could be applied. The disclosure should focus on the aspects of the invention which contribute significantly to its success, or technical contribution.

AI algorithms should be disclosed to a sufficient level to allow the skilled person to carry out the invention. Here, analogies could be drawn with the field of, for example, cryptography or various other computer implemented inventions. It is rarely necessary to disclose the full details of each step, merely an outline in sufficient detail to allow the skilled person to implement the invention using their common general knowledge. It is likely that variations may be seen between worked instances of an invention resulting from any chosen implementation. However, as long as these fall within the scope of the claims, this should be sufficient to constitute an infringement.

For example, where an AI machine is trained using training data, it may be necessary for any selection attributes of that training data to be disclosed, in particular if these contribute to the success of the AI machine. While a skilled person, seeking to carry out the invention based on this guidance, may not end up with an identical AI machine to that developed by the inventor, it may allow the benefit of the invention to be obtained. We do not see any need to fully disclose a training dataset or model weights. However, in some examples, publicly available standardised datasets may be used, which may be designed specifically for the purpose of AI training, or may include other publicly available data, such as government statistics, sporting records or the like. In such cases, the training dataset may be explicitly identified.

The patent application should be drafted such that it could be enforced against third parties who wanted to use the same data to train an AI machine.

As another example, tools such as https://cloud.google.com/automl make it very likely that a single dataset may be used to quickly train a variety of different ML models. For inventions involving the application of AI/Machine learning, we therefore consider precise model details to be largely

irrelevant and that the patentee should be given the benefit of the doubt when determining whether the advantage of the invention may be achieved across the full breadth of the claims.

Qu. 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 could be the case, in particular where working AI models have an impact on day-to-day life. For example, referring to the recent allocation of exam results based in part on postcodes during the covid19 pandemic, disclosure of the principles behind the algorithm also revealed the inherent unfairness thereof. However, this does not appear to be a subject for patent law, and instead could for example be associated with offering a service to the public and/or as part of a government service. It would be necessary for this to be distinct from patent law as not all AI inventions will be exposed through patent applications, as discussed above.

Qu. 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?

We consider current inventive step requirements to be adequate to cover inventions incorporating machine learning and other AI inventions. Rather than change the inventive step threshold, we suggest that the emergence of AI changes the tools available to the skilled person. Just as the person skilled in the art is considered to be generally computer literate, as technology evolves the skilled person should be considered to be literate in the use of machine learning and AI. The inventive step threshold may effectively become higher because the skilled person has better tools available to them with which to perform inventions, but this should be a consequence of the skilled person having access to said tools, rather than through a change in the inventive step requirements.

With respect to the example in the consultation document of an AI used to understand, learn and review dosage regimes for different pharmaceutical products, we note that the advent of computers had a similar effect of bringing calculations and data analytics processes that would have previously been impossible into the realm of the possible. The inventive step threshold proved adequate then, and we consider it will remain so in future. In assessing whether an AI determined dosage regime should be considered inventive, the configuration of the AI, the input parameters and input data might all be considered, for example, when assessing inventive step.

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

This does not seem to be necessary at the present stage. In the event that AI machines independently become a dominant source of inventions, it is not clear how those inventions could be assessed. For example, can an AI machine be inventive at all? Or are AI inventions just the inevitable result of a logical process, based on the inputs to the system.

Either way, at present, we feel that the only expansion of the concept of the person skilled in the art should be to "a person skilled in the art and having access to AI".

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

A very wide range of situations can be envisaged in the context of infringement of a patent by or involving AI. However, we suggest that the question of who is liable can sensibly be assessed with reference to existing practice. That is, if an AI machine performs a process, then the owner of that machine is responsible for the process and for its results. The consultation document highlights that there may be a number of humans behind infringement by AI including the developer, the

manufacturer, the owner or the product end user, and notes that the human behind an AI may well be unaware that they are infringing a patent. The consultation document also notes that patent holders may prefer an AI developer or manufacturer to be held liable, rather than the end user. We suggest that these issues are largely analogous to existing questions over infringement by non-AI processes and products. Patent holders and their advisors already target patent protection towards developers and manufacturers rather than end users. We note that in this context, the idea of AI as a tool for performing a task is a useful analogy. If the AI infringes a patent when operating as provided to the end user by the developer, it should be possible, assuming one's patent is suitably drafted, to choose to pursue the developer for sale of the tool, rather than the end user for use of the tool.

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

Difficulties in proving infringement by AI under certain circumstances are to be expected. However, questions around infringement detection and the feasibility of patent enforcement are already considered by a majority of patent applicants when deciding whether or not to seek protection for an invention. It would appear reasonable to suppose that patent applicants in the AI field will also take such issues into account when determining priorities for patent filing. The availability of patent protection of commercially useful scope for AI inventions in general will also have a significant impact on the scale of this problem in the future.

We thank the Intellectual Property Office for considering our comments above.

Yours faithfully

Gemma Robin
Partner
On behalf of Haseltine Lake Kempner LLP