

IPO Artificial intelligence and intellectual property: call for views

RESPONSE OF THE CHARTERED INSTITUTE OF PATENT ATTORNEYS (CIPA)

OVERVIEW

The Chartered Institute of Patent Attorneys (CIPA) is the professional and examining body for patent attorneys in the UK, representing virtually all the 2,500 registered patent attorneys in the UK, whether in industry or in private practice. Total membership is over 4,000 and includes judges, barristers, trainee patent attorneys and other professionals with an interest in intellectual property. CIPA represents the views of the profession to policy makers at national, European and international level, with representatives sitting on a range of influential policy bodies and working groups in the UK and overseas.

CIPA notes the following statement from the Consultation:

“The government also wants to make the UK a global centre for AI and data-driven innovation. Its mission is to increase uptake of AI for the benefit of everyone in the UK. This includes ensuring AI technology works for the people and making sure the UK has the best environment for developing and using AI.”

In our response, we include two particular suggestions to support the above mission and demonstrate UK leadership in this field:

- (i) a statutory modification of the definition of “inventor” to better accommodate the increasing use of AI technology (see our answer to Patents, Question 3)
- (ii) a change in practice regarding the patentability of core AI technology to incentivise the UK development of such technology (see our answer to Patents, Question 7).

Please do not hesitate to contact us if you have any questions about this consultation response.

Lead contributors: Simon Davies (Altair IP), Rachel Free (CMS), Mike Jennings (AA Thornton), Scott Roberts (BT), Matt Hervey (Gowling WLG).

Contact: Simon Davies, Chair, CIPA Computer Technology Committee, srd@altairip.co.uk.

NOTE ON DEFINITION OF AI

The consultation notes there is no single agreed definition of artificial intelligence and quotes a government definition: “technologies with the ability to perform tasks that would otherwise require human intelligence, such as visual perception, speech recognition, and language translation.” Some of our comments herein may only be relevant to a subset of systems falling within this definition.

PATENTS

For ease of discussion, we identify 3 types of invention relating to AI systems:

- 1) an invention relating to how an AI system works, for example concerning the internal structure/operation of the AI system;
- 2) an invention involving the application of an existing AI system as a tool, for example when used for a particular application; and
- 3) an invention to which the output from an AI system makes a significant contribution.

We refer to these different types of AI invention in our response below as Type 1, Type 2, and Type 3 respectively.

Question 1

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

Answer: The role of the patent system in relation to AI technologies is generally the same as for other areas of technology: primarily encouraging investment in technological research and development resulting in the public benefit of new or improved products. That investment can be by innovative businesses themselves, or patents can support third party investment in those businesses. In exchange, the publication of patent applications helps to disseminate technical information to support further innovation by others.

Question 2

Can current AI systems devise inventions?

Answer: The output from an AI system may make a significant contribution to an invention.

(We avoid the direct question of whether an AI system can “invent” or “devise an invention”, which is still a matter of real debate, because such terminology might be taken to imply some level of consciousness of the AI system that goes beyond the present capabilities of AI systems).

Particularly:

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

Answer: Any AI system is an agent of its human operator and can therefore be reasonably regarded as a tool for a human inventor. For example, in the drug discovery field, AI analysis of published material may be used to identify potentially useful candidate drugs or further indications for existing drugs.

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

Answer: Any human can make a claim for inventorship of a given invention; however such a claim is only valid if the legal requirements for inventorship are satisfied.

An AI developer may be an inventor, e.g. for a Type 1 invention that improves the operation of the AI system itself (see our Definitions above).

It also seems possible for a person who constructs data sets for training an AI system (such as to identify candidate drugs or for mass throughput chemistry) to be an inventor, e.g. if they perform the selection and curation of a training dataset in a non-obvious way which enhances the ability of the AI system to perform the desired task. As a specific example, such an

invention might involve a non-obvious choice of features to extract from the raw training data for input into the AI system.

It is hard to see how the mere use by a human of an AI system, without anything further, would support a claim to inventorship by that human, but a new use of an existing AI system which gives a surprising technical effect might be a potential invention.

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

Answer: Yes – for Type 3 inventions, the contribution of an AI system to an invention may be such that human involvement, for example in setting up, operating and processing the output of the AI system, falls short of being “the actual deviser of the invention” according to the prevailing interpretation of section 7(3) of Patents Act (PA) 1977.

At present it is not possible to obtain a UK patent in the absence of a human being who is the “actual deviser of the invention”.

Question 3

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

Answer: We see no need to identify an AI system as an inventor.

Rather, we suggest patent law should allow patenting of an invention in the situation contemplated in our answer to Question 2(c), in particular in a case where human involvement in the invention falls short of being the “actual deviser of the invention” due to a significant contribution from an AI system. To address this situation, section 7(3) PA 1977 could be amended to specify that the person responsible for the output from the AI system which provides this contribution should be regarded as an inventor.

Note that once a human has been identified as an inventor in this manner, other issues such as ownership can then be determined in accordance with existing law and practice.

(This approach is somewhat analogous to the provision of section 9(3) of CDPA 1988 regarding the authorship of machine-generated works).

It is important that UK law remains consistent with international obligations and practice. We believe that the above suggestion satisfies this requirement, since the determination of inventorship is a matter of national law. In addition, many other jurisdictions do not have the “actual deviser” definition of section 7(3), hence their starting position is different from the UK. To date, approaches in different countries have generally resulted in consistent practice, but this may be challenged as Type 3 AI inventions become more common. It is important for applicants to maintain as much consistency as possible.

Please note the above suggestion for a change to section 7(3) PA 1977 is by way of principle and is not intended to be a proposal for specific statutory language. The exact wording would require further consideration and review to ensure that the consequences are fully understood, including confirmation there is no inconsistency with UK international commitments.

A variation on the above suggestion would be for the person responsible for the output from the AI system to be denoted as the first owner (rather than as inventor). However, this approach may lead to a situation where no inventor is defined, which is more likely to cause a problem with international law and practice.

An alternative view has been expressed in CIPA that we should defer any legislative change because at present there is always a human inventor (or it can be accepted that without a human inventor no patent can be obtained), and/or it would be better to wait for a more harmonised approach to be adopted. As noted above, CIPA strongly supports harmonisation in

general, and the proposed change to section 7(3) could in fact provide closer alignment with other jurisdictions which do not have the same “actual deviser” definition of an inventor. In addition, if an agreed international approach is subsequently forthcoming (which does not seem to be imminent), further change is not precluded to align with this international approach.

An advantage of the proposed change to section 7(3) is to provide greater legal certainty now in respect of the ownership and validity of a patent for an invention having a significant contribution from an AI system (see our answer to Question 4 below).

Question 4

If AI cannot be credited as inventor, will this discourage future inventions being protected by patents?

Answer: As in our answer to Question 3, we do not propose that AI is “credited” as an inventor. Rather we suggest that a person is regarded as an inventor if they are responsible for the output from an AI system which provides a significant contribution to an invention.

We see three problems under the current situation if an AI system makes a significant contribution to an invention.

(1) If there is no human who can be regarded as the “deviser of an invention” (because of the contribution of the AI system), then a patent cannot be granted since no human inventor is identified. This may, in turn, discourage further research and development using an AI system because the results are not protectable by patenting.

(2) Some parties may identify a person as inventor to provide a basis for a patenting even though there may be some uncertainty whether the contribution of this person satisfies the current legal requirements for inventorship of section 7(3). This could potentially cause a problem downstream, for example in relation to doubt over the ownership and validity of the patent.

(3) A proprietor may downplay in a patent the contribution of an AI system to an invention (to avoid suggestion that a human is not the actual deviser); however, this may then reduce the accuracy and value of information provided to the public by publication of the patent application.

Would this impact on innovation developed using AI?

Answer: Yes - a company might be concerned that using an AI system for research and development might prevent a patent being obtained for the results, as per problem (1) above, or that the validity of such a patent might be placed in doubt, as per problem (2) above. In these circumstances, the company might choose not to use the AI system for the research and development (or they might choose not to perform the research and development at all). For example, a pharmaceutical company may be reluctant to use an AI system for the development of a new drug if this might compromise the patent protection available for the new drug, and this reluctance may slow the pace and increase the cost of development.

Would there be an impact if inventions were kept confidential rather than made public through the patent system?

Answer: Yes – there will be a reduction in knowledge sharing. In addition, the use of confidentiality by itself may be impossible in some situations, e.g. if the resulting invention is inherently disclosed when a product is placed on the market, such that there would be no effective IP protection for the invention. Even where the use of confidentiality is feasible, it may still be less effective than using a patent, e.g. for a start-up seeking investment based on its IP.

Question 5

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

Answer: No - we do not see a moral case for recognising an AI system as an inventor in a patent (and this would not occur under our suggested approach in response to Question 3).

Question 6

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

Answer: We do not recommend identifying an AI system as an inventor. Instead, we refer to our answers to Questions 2 and 3 above. Once an identification of a human inventor has been performed, ownership would follow in accordance with existing legal frameworks.

(We note that an AI system has no legal personality and so cannot be the legal owner of an invention. Crediting an AI as an “inventor” would therefore not solve the problem of identifying first and subsequent ownership of the invention).

Question 7

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

Answer: Yes – we see problems in particular in relation to Type 1 and Type 3 inventions (see our Definitions above).

Type 1: current UK practice often leads to this type of invention being rejected as excluded subject matter, namely as a computer program as such. We think this approach is inconsistent with how AI is seen in the wider world outside patents – AI is regarded as an integral part of modern technology and is developed by people with high levels of technical training and expertise in order to solve complex technical problems. We think Type 1 inventions generally make a technical contribution and hence should be patentable based on current law and decisions. This situation might be addressed, for example, by providing updated guidelines.

(For completeness, we note that the EPO sometimes rejects Type 1 inventions as a different form of excluded subject matter, namely as a mathematical method as such. Again, we think Type 1 inventions generally make a technical contribution and hence should be patentable).

Type 3: see our answers to Questions 2-4 above.

Question 8

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

Answer: We think any discussion of the future development of AI technology is necessarily somewhat speculative and we make no further comment.

Question 9

How difficult is it to secure patent protection for AI inventions because of the list of excluded categories in UK law?

Answer: It can be difficult to obtain patent protection for Type 1 AI inventions – see our answer to Question 7 above.

It is also difficult to obtain patent protection for Type 2 AI inventions that use AI as a tool for a purpose deemed to be non-technical (such as predicting share prices). However, we do not see this as an AI-specific issue or as a cause for concern. There is increasing international acceptance that a technical purpose or contribution is key to patentability for subject matter that would otherwise be excluded from patent protection and we support the status quo in this respect.

Where should the line be drawn here to best stimulate AI innovation?

Answer: We recommend a change in approach to the treatment of Type 1 AI inventions, since these relate directly to core AI technology as per our answer to Question 7 above. Furthermore, we think this change does not have to involve amendment of the law itself. Thus at present, the IPO tends to reject Type 1 AI inventions for being a computer program as such. It has already been noted in the Gower review that the IPO applies the computer program exclusion more broadly than the EPO (although the law is the same). A narrower interpretation of this exclusion, more aligned with the EPO approach, would facilitate the patentability of Type 1 inventions.

As noted in our response to Question 7, the EPO has developed a relatively recent practice of rejecting some Type 1 AI inventions for being a mathematical method as such. See our answer to Question 7 for our view on this practice.

More generally, we emphasise that AI systems are part of technology, and hence enhancements of such systems, even if implemented using a computer program and/or mathematical algorithm, should be patentable under existing legal provisions as representing the development of technology (Article 27 WTO TRIPS).

We recommend maintaining the existing treatment of Type 2 AI inventions regarding technical and non-technical subject matter, thereby retaining alignment with the broader treatment of computer-implemented inventions. However, we emphasise that this existing approach should not prevent the patenting of inventions directed to a specific improvement of an AI system or the use of such an AI system, even if the ultimate objective of such improvement is to support an application having a non-technical purpose.

We also note that there is some uncertainty at present regarding the treatment of simulations and other inventions with a “virtual” technical effect, pending the outcome of G1/19 from the EPO Enlarged Board of Appeal. This uncertainty extends to various potential uses of AI systems. We believe that patentability should be maintained for inventions having such a “virtual” technical effect providing they are performed for a technical purpose.

Question 10

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

Answer: We think such ethical oversight is more properly the responsibility of regulators and they should have the relevant power necessary to provide such oversight (this would be separate from the patent system – for example analogous to the regulatory regime for medical devices).

We also note that ethical problems are often viewed as non-technical, for example in relation to the right to be forgotten, improving fairness, etc. This may reduce the incentive for the development of AI solutions in this area, although we emphasise that any underlying technological innovation should be patentable. For further discussion, see <https://information.patentepi.org/issue-4-2018/technical-problems-in-ai-inventions-in-the-light-of-the-guidelines-for-examination-in-the-epo.html>

Question 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?

Answer: It is generally feasible for a patent application for an invention involving AI (e.g. a Type 2 invention) to include enough detail to allow a skilled person to perform the invention. The specific information to be included depends on the particular nature and use of the AI system for the invention. In some cases, it may be more difficult for a patent application involving AI to

include appropriate detail for a skilled person to perform the claimed invention, but this is also the situation for patents relating to other complex technologies and is not specific to AI.

One aspect of inventions involving AI which may be different compared with other fields of computing is the potential desire of an applicant to submit a significant amount of data, for example, a large training set and/or a specification of machine connection weightings. We think there are various ways to address this issue, such as by the introduction of a deposit system (see our answer below).

Is it clear how much information would be sufficient for a skilled person to be able to work the invention?

Answer: Yes – for example, at a practical level, there is extensive academic literature relating to AI technology and this represents one benchmark for the level of information to allow a skilled person to understand of developments in this field. In addition, there is already substantial UK and European case-law relating to sufficiency, and this case-law is generally applicable to AI inventions for determining how much information should be included in a patent application to meet the requirement of sufficiency. There may be challenges for patent offices and courts when considering evidence of the capabilities of AI systems that were available to a skilled person at a specific point in time, but the legal framework for this is in place.

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?

Answer: If an AI invention utilises a standard off-the-shelf AI system (e.g. TensorFlow) that is readily available to a skilled person, a patent application might only need to reference this AI system. In contrast, if an AI invention utilises a bespoke or customised AI system, then more details about the AI system might need to be specified in a patent application. We think this is generally in line with how other computer-implemented inventions are handled, and we do not see any major uncertainty in this regard.

What are the consequences if the details of AI algorithms need to be disclosed?

Answer: The consequences are that the applicant has the ability to obtain a patent and the public receives information about the relevant AI algorithms. Of course, inventors have a choice as to whether their commercial interest is best served by using the patent system or by maintaining confidentially (e.g. as a trade secret) and can then act accordingly. Again, we think this is generally in line with how other computer-implemented inventions are handled, and we do not see any particular issue in this regard.

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?

Answer: Yes – such a need might exist in some cases, for example if an invention directly relates to the selection, curation or configuration of the training data. Likewise, model weights of an AI model might be used to clearly disclose how to achieve the desired functionality of the AI model. Again, we think this is generally in line with how other computer-implemented inventions are handled, and we do not see any particular issue in this regard.

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?

Answer: Yes, it is clear – see our answers above to the first two parts of Question 11.

In practical terms, there might be some benefit in having a deposit system available for training data and such-like. We emphasise that this would be to facilitate **voluntary** submission of such training data, where the applicant believes such data to be helpful for the obtaining a patent.

As an alternative to a formal deposit system, the IPO could facilitate the inclusion of relevant data in an application, for example by specifying an appropriate format. In addition, the IPO could make such data readily available to third parties in a usable format (rather than just say PDF). An appropriate provision should be made regarding filing additional data in the specified format vis-à-vis page fees; in particular, if the training data submission is very large, this should not lead to very high page fees being required. Harmonisation of any such practise and capabilities across different patent offices would be very desirable.

What would be the effect if competitors could use data to quickly train a different AI model?

Answer: If a patent application discloses training data used for a first AI model, and this training data could also be used to train a second AI model and thereby produce the same effect and benefits as obtained with the first AI model, it appears reasonable for the inventor to be allowed claims that are broad enough to cover both the first and second AI models. This determination of appropriate claim scope follows existing principles of the patent system and these can also be applied to AI systems.

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?

Answer: The content of a patent application should provide support for the breadth of claims. For example, a patent application might show that a given result can be achieved across multiple sets of standard test data. Another possibility would be to show that a given result can be achieved with a range of different designs for the AI model and/or for different training sets. In addition, a skilled person is able to test an AI invention across the breadth of the claims to see whether or not the claimed result is reliably achieved. Again, this determination of appropriate claim scope follows existing principles of the patent system and these can also be applied to AI inventions.

Additional Comments: we do recognise that there are some practical challenges in handling AI inventions within the patent system. By way of example, some AI systems may be trained using a stochastic (random) process which may produce a variety of results, and only a small proportion of AI systems trained in this manner might provide a desired (positive) outcome.

We think such issues can generally be addressed using standard criteria to assess appropriate claim scope based on support, sufficiency, and range of inventive effect. In some cases, it may not be straightforward to perform such an assessment, but this is again true in some other fields of technology.

Question 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?

Answer: Yes – there could be such reasons outside patent law, for example in a regulatory context. However, akin to our answer to Question 10, we see such reasons as forming part of a regulatory environment separate from the IP system and they should not lead to placing any additional requirement on the content of a patent application itself.

Question 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?

Answer: Inventive step is determined with reference to a person skilled in the art. In line with existing practice, a person skilled in the art would be expected to use readily available tools and an AI system would be such a tool. The capabilities of this AI system and the nature of its use (and hence the outcomes that might be expected) should be assessed by looking at standard practice in the relevant field. Accordingly, we think AI can be accommodated within the current approach to (and level of) inventive step.

See also our answer to Question 11 in relation to claim scope.

Question 14

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

Answer: We think the person skilled in the art should remain a human. As noted in our answer to Question 13, such a person would be expected to use relevant tools, an AI system being an example of such a tool. The capabilities of this AI system and the nature of its use (and hence the outcomes that might be expected) would be assessed by looking at standard practice in the relevant field.

Question 15

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

Answer: We think this issue is part of a broader question regarding AI and the general law of tort and should be handled at this more general level (rather than at the specific level of the IPR system).

Question 16

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

Answer: There is a wide range of inventions involving AI. It may be difficult in practice to show patent infringement in respect of **some** AI inventions, depending upon the precise nature and use of any given AI invention.

This same issue exists in many other fields of computing (e.g. remote server operation; low level code processing), as well as in other areas of technology (e.g. manufacturing processes). Industry has already developed approaches to mitigate this problem, such as drafting claims with specific regard to the ability to show infringement and the use of legal procedures such as disclosure.

We don't have an estimate of the size of the problem, but it will not be relevant for all inventions involving AI. For those where it may be relevant, we regard the impact as manageable (in the same way it is for other areas of computing).

COPYRIGHT

Question 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 in the consultation document, however, we think the use by an AI system of third party copyright materials is more complex than indicated in the consultation.

As an example, we consider an AI system which is being trained to produce an image in the style of a given artist. The AI system may first download from the web other images made by that artist and extract features from these images, e.g. (convolutional) filtered versions of the images, histograms of colour palettes, and so on. The extracted features are then used to train the AI system, whereby the internal weights/connections of the AI system are updated to learn and represent the style of the given artist.

If we consider that making an image available over the web implicitly authorises downloading and viewing of the image, then (depending upon implementation) training the AI system may not involve making any further copy of the downloaded image (beyond the original download). It is unclear whether feature extraction would represent copying of a substantial portion of the image (this may depend on the particular selection of features). It seems less likely that updating the internal weights/connections of the AI system would represent a substantial portion of the downloaded image.

In summary, this seems to be a complex situation and greater clarity would be helpful – e.g. in relation to what represents a substantial portion of a copyright work in the context of feature extraction.

Question 2

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

We think in the case where an AI infringes copyright, it would typically be the operator who is liable. We recommend that such infringement be considered under the general law of tort, rather than creating some new provision specific to IP.

Question 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.

We think consideration should be given to providing greater certainty, e.g. in relation to the situation discussed in our answer to Question 1. This may result in a clarification of existing legislation or the creation of new exceptions.

Question 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 in our answer to Question 1, there is a need to provide additional clarity on how the existing legislation is applied in the context of AI systems. The answer to whether additional protection is needed may depend on the result of such clarification.

Question 5

Should content generated by artificial intelligence be eligible for protection by copyright or related rights?

We think the existing provision for computer-generated works according to section 9(3) CDPA 1988 should (and indeed does) apply to content generated by an AI system.

Question 6

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

This protection would be in line with current legislation in view of our answer to Question 5.

Question 7

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

We do not have any other specific issues to raise.

Question 8

Does copyright provide adequate protection for software which implements artificial intelligence?

We think the existing copyright system, in conjunction with other IP rights such as patents and trade secrets, provides suitable protection for software which implements artificial intelligence.

Question 9

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

We are not aware of any such obstacles. We note that various AI software is readily available under an open source license – e.g. Tensorflow is available under the Apache license which is generally considered to be a permissive license with relatively few restrictions.

DESIGNS

Question 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

Question 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?

We think an AI-created design should clearly be treated as generated by a computer and hence the author would be determined in accordance with s2(4) RDA and s214(2) CDPA.

Question 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?

We think the identification of the author according to the current legislation is appropriate and at present we are not aware of any reason to make a change.

We also note that contractual provisions between parties are already widely used to control ownership of designs (and IP in general), for example when commissioning a design, and such contractual provisions can also be used in relation to designs created by AI systems where so desired.

Question 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 – we see no benefit in this.

Question 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?

N/A given our answer to Question 4.

Question 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?

We think this could lead to uncertainty regarding the registerability and then subsequent validity of a design where this is no human authorship and also regarding the ownership of such a design. This uncertainty would apply to both computer-generated designs and also to AI-generated designs (which we see as a subset of the former).

Question 7

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

We are not aware of any other issues.

Question 8

Can the actions of AI infringe a registered or unregistered design? Can AI do the acts set out in law?

By way of example, an on-line ordering system might use an AI system to develop and control a user interface (UI) provided to users. For example, the AI system may generate slight modifications of the UI and adopt those modifications for which customers are found to place orders more quickly. Such an on-line ordering system might also use the AI system to create a design for a particular product (e.g. a teapot) and then to advertise and sell on-line the product according to this design.

In the above example, the modified UI or product design generated by the AI system could potentially fall within the scope of a registered design. If the actions of the AI system had been performed by a person, that person would infringe the registered design. However, since the AI system is not itself a legal entity, we do not think the AI system per se could infringe the registered design – rather the infringer might typically be the operator of the AI system. We think such infringement should be considered under the general law of tort, rather than creating some new provision specific to IP.

Question 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?

As long as a significant proportion of purchasers are still human for a particular product or design (which we think is likely for the foreseeable future), it seems appropriate to retain the existing human-based standard for the “informed user”.

Question 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?

As per our answer to question 8, we think this would typically be the operator, depending on the particular circumstances. We recommend that such infringement be considered under the general law of tort, rather than creating some new provision specific to IP.

TRADE MARKS

Question 1

If AI technology becomes a primary purchaser of products, what impact could this have on trade mark law?

By way of example, an AI system might be used to select and buy gifts for a particular recipient by searching the web to locate product reviews, pricing, functionality, and some form of matching between the ‘personality’ of the product brand and the intended recipient.

A competent AI system should be less likely to be confused than a human – e.g. there is no imperfect recollection, an AI system can be sensitive to small changes at the end of a word, unlike a human, and so on.

It is not impossible that a less competent AI system might be more easily confused in some respects than a human. For example, a well-known brand might be represented by a particular image of a dog. A third party might deliberately market itself using another image of a dog which to a human is quite different from the brand image, however the AI system might just treat them both as dogs and so confuse the two brands. This might be better regarded as a shortcoming of the AI system rather than an issue under trade mark law.

Question 2

Are there or could there be any difficulties with applying the existing legal concepts in trade mark law to AI technology?

Yes, there might potentially be some difficulties – see our answers to the questions below.

Question 3

Does AI affect the concept of the “average consumer” in measuring likelihood of confusion?

As per question 1, it is expected that an AI system will have a different (hopefully reduced) likelihood of confusion compared to an average human consumer. In theory this might support the co-existence of marks that currently would be regarded as confusingly similar.

However, as long as a significant proportion of purchasers are still human for a particular product or brand (which we think is likely for the foreseeable future), it seems appropriate to retain the existing human standard for the “average consumer”.

Question 4

What is the impact of AI on the drafting of section 10? Can AI “use in the course of business” a sign which may be confusingly similar or identical to a trade mark?

Yes – see our answer to Question 5 below for an example of an AI system using a sign which might be confusingly similar to a trade mark. Whether this is “in the course of business” would generally have to be determined with reference to a legal person (human or company) responsible for operating the AI system.

Question 5

Can the actions of AI infringe a trade mark?

The actions of an AI system could potentially lead to infringement of a trade mark. For example, an AI system may generate different versions of a slogan to advertise a product, the AI system being trained to enhance the slogan based on those versions which produce the best response. This training may lead the AI system to create and use (e.g. by placing adverts on the Internet) a slogan which ends up being confusing similar to an existing third party trade mark.

Since the AI system is not a legal entity, it cannot be liable itself for the trade mark infringement, rather this liability would fall on some legal person (human or company) responsible for operating the AI system. We think this should be considered under the general law of tort rather than specific provisions for IP.

Question 6

If AI can cause trade mark infringement, does this shift who could be liable? Should it be the owner, the operator, the programmer, the trainer, the provider of training data, or some other party?

As per our answer to question 5, an AI system is not a legal entity, so cannot be liable itself for trade mark infringement. Rather, liability would fall on some legal person (human or company) responsible for operating the AI system. We think this should be considered under the general law of tort rather than specific provisions for IP and may be dependent on particular details of a given AI system (including how it was developed and trained). It is also possible that contractual provisions between the different parties such as the supplier and the operator may be relevant.

It seems conceivable that “in the course of business” might be determined with respect to one party while the ultimate liability for any trade mark infringement might fall (at least in part) on a different party.

TRADE SECRETS

Question 1

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

Yes. The importance depends on various factors, such as:

- *no patent protection is available for legal reasons – e.g. non-technical/excluded subject matter

- *it is difficult to detect infringement or otherwise enforce the patent - e.g. a potentially infringing system is likely to be inaccessible behind a cloud server or sited offshore in a remote jurisdiction

- *a lot of effort for developing the AI system relates to aspects which might not be inventive per se, e.g. transforming a training data set previously used with an older AI system from an old format to an updated format for input to a new AI system

The presence of the factors in the above following examples would tend to increase the importance of trade secrets. Conversely, trade secrets might have little or no role if the AI system is used to create a product which is marketed to consumers and so accessible for public inspection.

In addition, AI systems are often considered to have a “black box” nature, i.e. it can be difficult to understand how or why a specific input produces a specific output. Trade secrets in the form of user experience and know-how may be important in obtaining the best performance from such systems.

Question 2

Does AI impact UK trade secret law? Does UK trade secret law give adequate protection to aspects of AI technology where no other intellectual property rights are available?

We are not aware of any direct impact of AI on UK trade secret law (rather than vice versa). Possibly the nature and form of AI systems may make it harder to identify specific information which is protected as a trade secret, but in general, we think UK trade secret law does give adequate protection to AI technology.

We see different forms of IP protection as complementary rather than necessarily as alternatives. Thus trade secrets might be used in conjunction with copyright and/or patents to protect an AI system – this is already common practice in the computing sector.

Question 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 of using trade secrets include:

- *no cost

- *immediately and automatically acquired

- *available for items that are not suited to patent protection, e.g. non-technical material such as a collection of training data relating to stock market prices

- *good for protecting a collection of details which individually are not worthy of patent protection, but collectively can bestow competitive advantage

- *no public disclosure to competitors (as would happen with an 18 month patent publication)

Disadvantages of using trade secrets include:

- *not suitable if commercialisation necessarily involves disclosure to the public (such as by marketing a product)

*licensing and supporting investment may be harder than for registered rights (since the scope of the trade secrets is less certain and potentially more fragile)

*not effective against third party creation

*limited effectiveness in relation to employees leaving with knowledge

*may be harder to obtain public recognition of technological expertise

Note that the above advantages/disadvantages apply to computing technologies in general and also across other industry sectors. Thus the situation with AI is not very different from existing computer technologies.

Many AI systems (and also many computing systems in general) are likely to rely on at least some trade secrets protection, and this may be supplemented as appropriate by the use of patent and/or design protection according to the particular circumstances.

In terms of not sharing information, this could potentially hinder AI development because third parties are excluded from knowledge that could help them develop their own AI system. On the other hand, not sharing information (e.g. with potential competitors) could help to protect investment in the development of AI systems, and therefore encourage greater investment and development.

We think the existing IP system provides a reasonable balance between the above considerations for the computing industry in general, and this balance should apply also to AI development.

Question 4

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

Potentially yes. For example, if an AI system is used to screen people, e.g. as part of a credit check or a job selection process, there is a concern the AI system might incorporate a subtle bias, e.g. relating to gender, ethnicity, and so on. If the details of the AI system are subject to trade secret protection, it may be more difficult for users and other interested parties to determine whether or not the system is free from bias. Another example might be in trying to assess the safety of an AI system used to control a driverless car - if the details of the AI system are subject to trade secret protection, it may be more difficult to make a reliable assessment.