Scott Roberts
European Patent Attorney

British Telecommunications plc
Intellectual Property Department
Association Number 854
1 Knightrider Street
London EC4V 5BT

Email: scott.4.roberts@bt.com
Telephone: +44 (0) 33 162 049 69
Facsimile: +44 (0) 20 760 068 91

Intellectual Property Office
Concept House
Cardiff Road
Newport
United Kingdom

BY EMAIL TO AIcallforviews@ipo.gov.uk

27 November 2020

Re: Artificial intelligence and intellectual property: call for views

Response of BRITISH TELECOMMUNICATIONS plc

Summary

- BT recognises and supports the government's goal of making the UK a global centre for artificial intelligence (AI) and data-driven innovation by increasing uptake of AI for the benefit of everyone in the UK.

- Innovators depend on intellectual property rights (IP) to protect and commercialise their innovations in the field of AI. In particular, patents incentivise investment in research and development (R&D) by protecting inventions as a property right while at the same time publishing the details of inventions for the wider good.

- Inventions in the field of AI should be susceptible to patent protection where the criteria for patentability, including novelty and inventive step, are satisfied. It is conceivable that an AI system may assist in the discovery of novel ideas in a technical field, some of which may be non-obvious in that field applying existing standards for inventive step. These ideas should also be eligible for patent protection.

- AI is a tool and does not devise inventions. An AI tool always operates in conjunction with human input, control and/or guidance. However, the current

criteria in patent law for determining inventorship for a novel idea do not make it clear that a person making arrangements necessary for an invention to be devised with the assistance of AI should qualify as an inventor. It is therefore proposed to clarify the definition of "inventor" in patent law in the case of computer-assisted inventions such as by including a person by whom the arrangements necessary for devising an invention are undertaken.

- In view of the purpose of identifying an inventor for a patent application to determine who has the right to be granted a patent, an AI or computer system should never be designated as an inventor or co-inventor. With an appropriate clarification to the criteria for inventorship, the existing provisions in respect of the right to file a patent application are satisfactory. These provisions ultimately identify the owner of a patent application and any resulting patent.

- A new idea discovered with the assistance of AI is suitably assessed for inventive step applying the existing standard defined in patent law and applied in the case law. It is not necessary to extend the legal fiction of a "person skilled in the art" to include "a machine trained in the art". Such an extension would serve only to expand the fictional capability of the person skilled in the art to have recourse to trained machines (i.e. AI). A determination of whether a notional skilled person may or may not have recourse to AI is part of an analysis of inventive step for an invention and is a question to be addressed in the context of the invention, the relevant state of the art and the common general knowledge of the skilled person.

- An infringer of a patent is strictly liable under tort law. That a legal person performs an infringing act through the agency of an AI does not affect their liability.

Introduction

The UK took its place at the forefront of the digital revolution during the latter half of the twentieth century. The UK's disproportionately high intellectual contribution to R&D in the computer and communications sciences led to early developments in theories of computation, computer processing systems, the software industry, telecommunications, the microcomputer and home computer revolutions and high speed network access. The UK continues to take its place as an international thought leader based on a combination of a highly skilled workforce and quality research and development, positioning the UK once again at the forefront of a forthcoming *4th industrial revolution* (4IR).

AI is one of the building blocks of the 4IR, along with technologies including Internet-of-Things, distributed transactional databases, virtual and augmented reality and additive manufacturing. These enabling technologies have the capacity to transform industry, business and the way we live and work. The opportunity for the UK is to lead the world in 4IR expertise, technology, products and services which, in turn, will drive growth and international competitiveness. The pathway to this opportunity is to grow R&D, skills and experience nationwide.

The products and enabling technologies of 4IR are predominantly intangible, including R&D outputs, designs, software, data and ideas. Innovators depend on IP rights to protect and commercialise these intangibles. Patents incentivise investment

in innovation by protecting inventions as a property right while at the same time publishing the details of inventions for the wider good. Without patents the incentive to innovate is diminished, protection for inventions is reduced to trade secrets, and knowledge sharing is curtailed. Furthermore, the opportunity to commercialise innovations can be affected as ideas fall available for anyone to use, even those arising as the product of considerable research investment.

We recognise and support the government's goal of making the UK a global centre for AI and data-driven innovation by increasing uptake of AI for the benefit of everyone in the UK. The present consultation invites views on the implications of AI technology for IP policy and in this response we focus on patents. The consultation is explicit that speculation of concepts of AI superintelligence (so-called "strong AI") is not occasioned, and we agree that the philosophical discourse in relation to strong AI is not aligned with the current and foreseeable state of the art. In this regard, references to AI in this response are references to the current and foreseeable state of AI technologies including, inter alia, machine learning, evolutionary computation, fuzzy systems and probabilistic methods. That is not to suggest that significant and unforeseeable developments will be made in these and other AI techniques, though there is no evidence to suggest that strong AI will be forthcoming in the foreseeable future.

Inventions in the field of AI arise in three main respects:

   I. Inventions created in the furtherance of the field of AI, such as developments and improvements to existing AI techniques or new AI techniques that advance the technical field.
   II. Inventions involving the application of AI to address a problem in a technical field, such as speech synthesis, image recognition, or network intrusion detection.
   III. Inventions including novel ideas discovered with the assistance of AI, such as novel outputs of a trained machine learning algorithm.

All such inventions should be susceptible to patent protection where the criteria for patentability, including novelty and inventive step, are satisfied.

The specific questions raised in respect of patents in the consultation are addressed in turn below.

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

> As outlined above, the patent system primarily encourages investment in technological research and development. New ideas in the field of AI can be wide-ranging and broadly applicable and no other IP right is suitable for protecting such innovations across their breadth.
>
> While the ability for innovators to protect their intellectual property through patents is essential to drive growth in R&D investment, it also brings important and valuable additional benefits. An innovative ecosystem is fundamental to an internationally competitive economy and grows demand for a high-quality and highly skilled workforce. Incentivising R&D investment, such as through patents, therefore drives demand for, and improvements to, world-class

education, world-beating universities, high quality professional services and world-class legal and policy frameworks. The UK can be at the centre of the 4IR and can lead internationally while reaping benefits across all regions as one more way to level-up the country.

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

In principle, it is conceivable that an AI system may assist in the discovery of novel ideas in a technical field, some of which may be non-obvious in that field applying existing standards for inventive step. Examples of such new ideas discovered with the assistance of AI systems are indicated by the contribution of Beat Weibel to WIPO meeting WIPO/IP/AI/GE/19[1] and the DABUS[2] applications.

Notwithstanding this possibility of an AI assisting the discovery of novel ideas, AI is a tool (or a suite of tools) and does not devise inventions. An AI tool is the creation of, configured by, operates under the instruction of, and with objectives directly or indirectly specified by, a person. An AI is a type of conceptual computer system, the programming of which is performed by training, supervising, defining objectives and/or fitness functions, defining adaptation functions and the like.

The reasons for requiring the identification of inventor(s) of an invention in a patent application are twofold: 1) to enable the right of attribution provided in, inter alia, Article 4ter Paris Convention; and 2) to determine ownership of resulting property deriving from the right to apply for a patent. In respect of ideas discovered by AI, both purposes are irrelevant: there is no right of attribution for a computer system such as an AI tool; and there can be no proprietary ownership by an AI tool. It follows that an AI tool cannot be an inventor at least because: the law does not permit it to be[3]; and the principal purpose of determining inventorship is to determine ownership, and the law does not permit AI systems to own property.

Some stakeholders identify ideas discovered, at least in part, by or with the assistance of AI. Where a stakeholder considers they are unable to identify a human inventor for an idea, the idea is essentially excluded from patentability since the patent system requires the identification of a human inventor[4]. An inability to identify a person satisfying the criteria of "inventor" for an invention places the invention outside the realm of the current patent system.

---

[1] Day 1 (afternoon) at 02:05:49
[2] UK patent applications GB1816909.4 and GB1818161.0 purportedly conceived by the "creativity machine" DABUS, details of operations of which are described in US patent publications: 5,659,666; 7,454,388 B2; and 2015/0379394 A1
[3] Thaler v Comptroller General [2020] EWHC 2412 (Pat) at 40
[4] *Ibid.*

There is, therefore, a clear need to bridge a gap between the provisions on inventorship in law and the determinations of stakeholders in respect of inventorship of novel AI discoveries. Noting that:

    i.   the Patents Act specifically requires that the "inventor" is the "actual deviser" of an invention[5];

    ii.   the right to be granted a patent is specifically tied to inventorship[6], and to rely on "a subjective belief on the part of an applicant that that applicant is entitled to apply for a patent" is considered by the courts to be "nonsense"[7];

    iii.   a prevailing standard for identifying a person as an inventor is that, on the balance of probabilities, the person had made a relevant contribution to one of the inventive concepts in a patent[8];

    iv.   it is not sufficient, for inventorship, for a person to merely contribute to features of a claim[9]; and

    v.   for patentability, an invention must involve an inventive step[10] meaning it is not obvious to a *person* skilled in the art[11],

the gap is identified by any applicant who, in applying the law, is unable to identify a human inventor. It is essential that this gap is bridged in order that patent protection is available for new ideas discovered with the assistance of AI. Failure to bridge this gap can lead to an erosion of the incentive for investment in AI research and development and a reduced disclosure of AI research through published patent applications as stakeholders turn to trade secrets to protect their novel ideas. It is proposed that this gap is bridged by providing explicit provisions for determining the inventor of inventions devised with the assistance of a computer through a change to the definition of "inventor" in the Patents Act.

For an invention including a novel idea discovered with the assistance of AI, the problem is how to attribute inventorship to a person where the person does not satisfy the existing requirements for such. Parallels can be drawn with UK copyright law and the protection afforded for computer-generated works, and the characterisation of such a person in the Copyright, Designs and Patents Act at S9(3) is a promising start: "…the person by whom the arrangements necessary for…" In respect of inventions, the relationship between such person and the invention requires careful definition to ensure the right to be granted a patent rests with the appropriate person. In extremis, challenges can include identifying the owner of novel ideas discovered with the assistance of AI developed by a first party, hosted by a second party, having a configuration (such as training data) provided by a third party, executed under the instruction of a fourth party, addressing a problem specified by a fifth party and/or having application to a problem recognised by a sixth party. While such challenges may be addressed contractually, it is not sufficient for stakeholders to depend exclusively on such instruments for these important determinations.

---

[5] Section 7(3) UKPA

[6] Section 7(2)(a), (b) and (c) UKPA

[7] Thaler v Comptroller General [2020] EWHC 2412 (Pat), at 29

[8] University of Southampton's Applications [2004] EWHC 2107 (Pat); [2005] R.P.C. 11

[9] Henry Brothers (Magherafelt) Ltd v Ministry of Defence and Northern Ireland Office [1997] R.P.C. 693, Pat Ct; [1999] R.P.C. 442, CA

[10] Section 1(1)(b) UKPA

[11] Section 3 UKPA

Accordingly, the wording of any proposed provision in patent law to address inventorship for novel ideas discovered with the assistance of AI is of fundamental importance. For an author of a computer-generated copyright work, this is simply recited as "…the person by whom the arrangements necessary for <u>creation</u> of the work…" However, given the range of persons potentially involved in the discovery of a new idea by an AI, such an approach is not suitable for inventions. Furthermore, even references to persons recognising a "solution to a problem" are inadequate, as inventions can reside in the identification of a new and non-obvious problem itself.

It is appropriate to build upon the existing criteria and standards for determining inventorship starting with the requirement of S7(3) UKPA that "inventor" means "the actual deviser of the invention" and that S7(2) and (3) UKPA is "an exhaustive code for determining who is entitled to the grant of a patent"[12]. This requires first an identification of the inventive concept (the "invention") and subsequently an identification of "the natural person who came up with the inventive concept"[13].

Some stakeholders argue that a person asking an AI to solve a given problem, providing only as input the state of the art, is insufficient to render the person an inventor. Such analysis is incorrect and based on a false premise because the requirement to identify an inventor necessarily requires the existence of an invention. If a machine is able to solve a given problem so readily based only on the provided state of the art and presentation of a problem, it is unlikely that the solution could be considered inventive. This is especially the case if the AI is a commoditised or known AI, and therefore readily available to the person skilled in the art. Five possibilities can be contemplated where the solution may involve an inventive step, each of which emphasises the involvement of a human inventor:

- Firstly, if the AI is itself novel and non-obvious, in which case the developer of the AI has created an invention in the AI and the products of the AI;
- Secondly, if the use of an AI to address the problem is not obvious, in which case the person turning to such AI has invented;
- Thirdly, if the problem itself is novel and not obvious in which case the person identifying the problem which is solved with the assistance of the AI has invented;
- Fourthly, if the input data was carefully selected so as to lead the AI to discover a novel and non-obvious solution to the problem, in which case the person making this special selection of the input data has created an invention with the assistance of the AI by specifically instructing the AI on the basis of the special selection; and
- Fifthly, if a particular output of the AI is selected by a person who recognises its usefulness and efficacy in addressing the problem, in which case the person making the selection based on this recognition has invented with the assistance of the AI.

A fundamental premise is that AI is specifically tied to a goal or motive which can be specified across a number of dimensions including, inter alia: the

---

[12] Yeda Research and Development Company Ltd v. Rhone-Poulenc Rorer International Holdings, [2007] UKHL 43, at 20
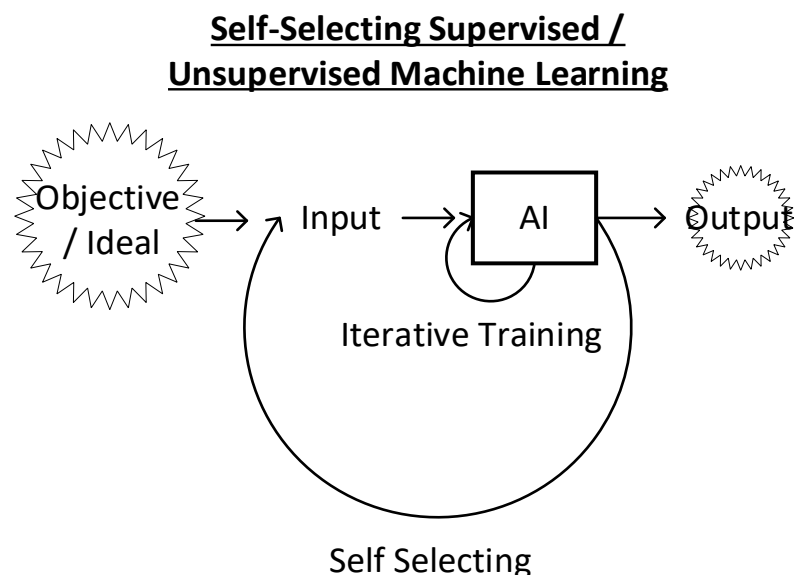[13] *Ibid.*

nature of the input data; the selection of an AI algorithm; the nature and arrangement of the AI algorithm; and the application of the output. Academics suggest that, even as AI *strengthens* towards "strong" AI, goal setting and motivation remain essential[14].

It is therefore worthwhile considering the use of AI in other contexts and how human involvement is inherent. First considering a conventional machine learning approach in which an input training data set is used to train an AI including a machine learning algorithm which generates some output such as a classification, an optimisation, a discovery or somesuch:

Supervised / Unsupervised
Machine Learning

Input ⟶ AI ⟶ Output

Iterative Training
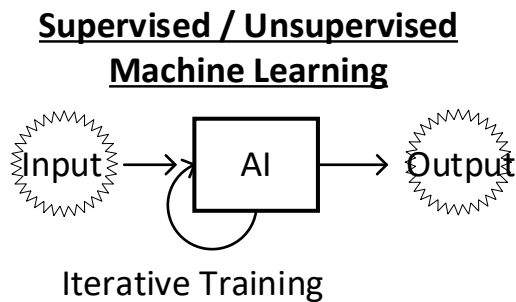
Human involvement is indicated by the pointed circles and extends at least to the selection of input data and/or the recognition of the application of the output data, and may further extend to the nature, operation and configuration of the AI itself, novelty of the problem addressed, and other conceivable aspects of novelty.

Conceivably, such an AI may find, generate or select its own input so becoming essentially a self-selecting machine learning algorithm:

Self-Selecting Supervised /
Unsupervised Machine Learning

Objective / Ideal ⟶ Input ⟶ AI ⟶ Output

Iterative Training

Self Selecting

---

[14] "General purpose intelligence: arguing the orthogonality thesis", Stuart Armstrong, January 2013, Analysis and Metaphysics 12:68-84; and the Reimann Hypothesis considered in "The Ethics and Risks of Developing Artificial Intelligence", Russell, Stuart J. & Norvig, Peter, in "Artificial Intelligence: A Modern Approach", 2003, Prentice Hall. ISBN 978-0137903955.
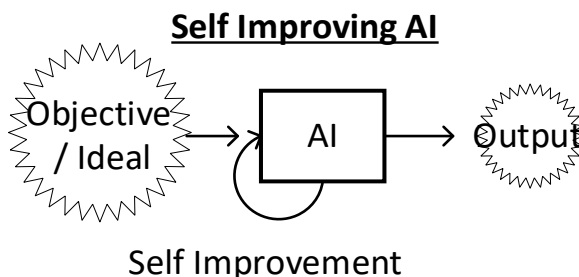
Such an algorithm includes human involvement including the determination, definition, selection or specification of one or more objectives, ideals or fitness criteria/functions for the selection of the input data.

Evolutionary algorithms are known to evolve and adapt using techniques including reproduction, mutation, recombination, and selection:

Evolutionary AI

Objective / Ideal → AI → Output

Improvement (incl. stochastic)
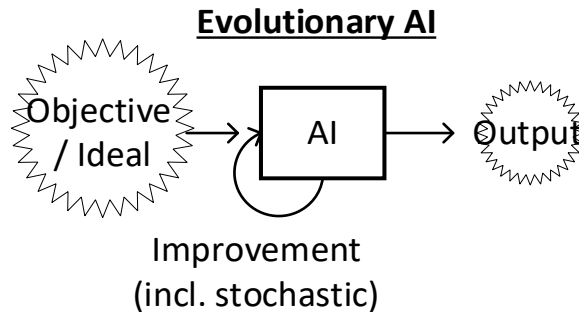
These algorithms include human involvement in the definition of the objective, ideal or fitness function and/or the identification of the motive/incentive for evolution. Even where such algorithms employ stochastic methods there is an objectively determinable efficacious subset in each generation, with efficacy being human defined.

Some AI can be described as self-improving such as self-adapting, reprogramming or adjusting AI:

Self Improving AI

Objective / Ideal → AI → Output

Self Improvement

These algorithms depend on human involvement including the objective, ideal or fitness function driving such improvement and/or the identification of a motive/incentive for adjustment.

Even where multiple of these AI techniques are combined or daisy-chained in a way that may give an impression of a completely autonomous system, it can be seen that human involvement is a requisite in at least one aspect of the AI. Where such an AI is involved in the discovery of a novel idea, the identification of human involvement is a suitable basis for recognising human contribution to the inventive process.

Recognising these exemplary identifications of human inventions for broad illustrative applications of AI assisting in the discovery of novel ideas, a reasonable legislative provision for identifying an inventor in the case of an invention devised with the assistance of a computer may be that "the inventor shall be taken to be the person(s) by whom the arrangements necessary for devising the invention are undertaken."

It is to be recognised that, however one views the question of inventorship for new ideas discovered with the assistance of AI, the challenge of determining the inventor considering the potentially many parties involved remains. Even stakeholders arguing that there is no gap between the legal provisions on inventorship and the state of technology must contemplate which of the potentially numerous parties involved with an AI is entitled to be granted patent applications. A failure to clarify this issue (such as by legislation) renders the question unanswered, and potentially unanswerable, for cases where contractual arrangements for ownership are absent or inadequate.

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

No. Whereas AI can assist in the discovery of novel ideas, AI cannot devise inventions. An AI tool is the creation of, configured by, operates under the instruction of, and with objectives directly or indirectly specified by, a person. An AI is a type of conceptual computer system, the programming of which is performed by training, supervising, defining objectives and/or fitness functions, defining adaptation functions and the like. Irrespective of its lack of legal personality, its inability to own or assign property and its exclusion from many moral considerations, AI anyway lacks motive, purpose, goal or object save for that which is programmed into it whether directly or indirectly by a person. Even an AI that discovers novel ideas through serendipity or stochastics does so on the basis of an algorithm and input data set provided ultimately by a person. Even if that serendipity is guided by goals, such goals are provided ultimately by a person. It is the person that is the inventor.

Furthermore, there is no purpose or need to identify an AI as an inventor. As outlined above, the reasons for requiring the identification of the inventor(s) of an invention in a patent application are twofold: 1) to enable the right of attribution provided in, inter alia, Article 4ter Paris Convention; and 2) to determine ownership of resulting property deriving from the right to apply for a patent. In respect of ideas discovered by AI, both purposes are irrelevant. There is no right of attribution for a computer system such as AI, and there can be no proprietary ownership by AI.

Some stakeholders advance a view that attributing a new idea discovered with the assistance of AI to a person diminishes the accomplishments of people who have legitimately created patentable works[15]. These views arise from the false premise that such ideas are not the legitimate conception of the person (see the response to Question 2 above). Such analysis is also circular based on the current state of the law that may be interpreted to preclude the acknowledgement of *a person making the arrangements necessary for devising an invention* as an inventor. The possibility of such an interpretation motivates clarification in the law, not a fundamental shift in the meaning of the term "inventor". Clarifying the law to recognise the human inventor for new ideas discovered with AI is a positive evolution of the provisions that specifically incentivises investment in AI research and results in ownership of inventions conceived with the use of AI to the appropriate natural or legal person.

---

[15] "The Reasonable Robot", Ryan Abbott, Cambridge University Press

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?

Being unable to credit AI as an inventor does not discourage future inventions being protected by AI. This is especially the case if the contribution of a human inventor is recognised and accommodated, such as by way of the legislative change considered above in response to Question 2.

To the extent that there is any disincentive to research in the field of AI and its applications, this arises due to a failure by some stakeholders to attribute their inventions to human inventors. Failing to identify a human inventor places an invention outside the realm of patent protection due to the requirement for a human inventor. Such failure may arise due to interpretation of the provisions for inventorship in patent law and this is readily addressed through the legislative clarification outlined in response to Question 2. It is a failure to make such legislative clarification that may result in AI inventions being kept as trade secrets.

Indeed, permitting designation of AI as inventor introduces considerable uncertainty as to the right to apply for a patent and ownership of resulting patents. We believe permitting designation of AI as inventor would itself act as a disincentive to pursuing patent protection, and therefore a disincentive for investment in AI R&D. An AI involved in assisting in the discovery of new ideas may exist in multiple capacities at the same time including in relation to: its owner; its lessee or tenant, such as in a cloud-computing context; its user or operator; a person responsible for or contributing to its configuration; a person recognising the value and contribution of its output; and others. Seeking to untangle these relationships from a property-ownership perspective is much less straightforward than seeking to recognise the persons involved in the devising of an invention. BT therefore believes the appropriate approach is to prioritise clarifying the human inventors involved in devising an invention.

Some stakeholders consider that "although the prospect of a patent would not directly motivate AI to invent, it will inspire the people who build, own, and use AI"[16]. Such views highlight the need to ensure new ideas discovered with the assistance of AI are susceptible to patent protection and also the link between people and the inventive process – i.e. the motivation, goal or objective. The need to provide patent protection is clear as an incentive to continued research, as is the dependence on a person as human inventor.

Some stakeholders argue that ownership of a patent including a novel idea discovered with the assistance of AI should rest with the owner of the AI[17]. Such an approach neglects to accommodate situations where the inventor is other than the owner of the AI. This is particularly likely in the field of AI since much AI is provided as software available from software service providers, licensors, cloud providers and the like. In contrast, an approach involving legislative change such as that outlined in response to Question 2 would serve

---

[16] *Ibid.*
[17] Thaler v Comptroller General [2020] EWHC 2412 (Pat), at 49(2)

to determine ownership on the basis of the person(s) by whom the arrangements necessary for devising an invention are undertaken.

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

There is no moral reason for recognising AI as an inventor. Some stakeholders advance a view that attributing a new idea discovered with the assistance of AI to a person diminishes the accomplishments of people who have "legitimately" created patentable works. This is incorrect and based on a false premise as explained above in response to question 3.

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

An AI should never be named as an inventor for a patent because AI does not devise inventions (see the responses to question 2 to 5 above). A human inventor should always be listed at least because all new ideas discovered by AI involve a human inventor. The identification of a particular human inventor can be aided by a legislative change such as that outlined in response to question 2. Ownership is suitably defined on the basis of inventorship as currently provided in the UKPA

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

Considering inventions across in the following three categories:

I.   *Inventions created in the furtherance of the field of AI*: such inventions can fall-foul of the excluded subject matter provisions codified in S1(2) UKPA and Article 52(1) European Patents Convention (EPC). In particular, improvements to AI algorithms and techniques without specific practical application can be excluded as mathematical methods or computer programs as such. We are generally comfortable with the interpretation of these provisions in the UK Patents Act and European Patent Convention, though note that inconsistencies can arise in the examination of patent applications between the UK Intellectual Property Office (UK IPO) and the European Patent Office (EPO) due to differing methods of applying the exclusions. Additionally, we recognise that developments in AI techniques constitute valuable contributions to the state of the art and their protection by patent without requiring recitation of a specific application would fairly reflect that contribution.

II. *Inventions involving the application of AI to address a problem in a technical field*: the application of AI to particular problem domains is normally readily protected by patent.

III. *Inventions including novel ideas discovered with the assistance of AI*: it is in the use of AI to assist in the discovery of new ideas that some patent applicants struggle to identify a human contribution sufficient, in their view, to warrant designation of an inventor. Such inventions should not be excluded from patent protection for only such reasons

and the clarification in respect of the human contribution to the devising of an invention described above is warranted.

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

Current and foreseeable AI technology will always involve a human inventor. An AI tool is the creation of, configured by, operates under the instruction of, and with objectives directly or indirectly specified by, a person. The prospect of the development of "strong AI" which may necessitate a review of the involvement of humans in the invention process is considered unlikely for the foreseeable future.

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?

Refer to the response to Question 7 above. The current approach to interpreting the excluded categories at the UK IPO and EPO are reasonable though some inconsistencies arise between the two. The outcome awaited in the currently pending referral to the EPO Enlarged Board of Appeal in G1/19 may have an impact on the patentability of some AI inventions. If this is the case, then further dialogue on the nature and scope of the exclusions will be required.

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

We consider that the answer to this question is no. Ethical oversight of AI includes questions of transparency, explainability and bias. While we recognise the importance of these issues, we believe it is not appropriate for them to be addressed through IP provisions such as patent law and practice.

Regulatory requirements to assure the social acceptability of AI should be considered separate to the patent system. While it is noted that patent applications must disclose an invention in a manner that is sufficiently clear and complete for it to be understood, this requirement serves to ensure only the sufficiency of disclosure of an invention in an application. Such disclosure need not extend to disclosure for the purpose of full transparency of an AI algorithm, or full explainability of an AI technique. Indeed, in some cases, a patent application for an invention devised with the assistance of AI may be directed to the output of the AI, in which case there may be no need to disclose the AI mechanism itself in order to sufficiently disclose the output itself. Accordingly, patent law is not an appropriate or effective home for these regulatory issues.

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?

We consider that the answer to this question is no. The disclosure of inventions including applications of AI or devised with the assistance of AI is no different to the disclosure of other computer implemented inventions. Methods, algorithms and flowcharts can be used to disclose the methodology

of an AI invention. There is unlikely to be a need to disclose training data or trained weights for machine learning algorithms or the like since inventions arise in the methodologies employed. Where novel outputs are discovered with the assistance of AI such as novel physical objects, novel chemical compositions and the like, disclosure of the output is sufficient. It is not necessary to disclose how the output was discovered to sufficiently disclose the output itself.

In particular, we do not believe the detailed disclosure of training data or trained neural network data structures (e.g. including trained inter-neuron weights) is necessary for a patent application, especially where a new idea is discovered based on the characteristics of the novel selection of training data, the output of a trained AI or the nature of an AI algorithm itself. Such data might only be beneficial as examples in a patent specification provided at the discretion of the applicant.

As a purely illustrative example, consider an AI algorithm for a self-driving vehicle applied in an *exemplary* novel way to differentiate a vehicle and a pedestrian on the road. Such an AI may be trained based on a large dataset of labelled training examples each including either a vehicle or pedestrian. The trained AI would then execute as a "black-box" to differentiate vehicles and pedestrians in-use. Were such an idea to be novel, a patent application to such an application of AI need not provide the large, labelled training dataset. A characterisation of the data set with examples and an indication of a volume of training data required would suffice.

In machine learning, the efficacy of a trained AI algorithm to achieve a certain goal is dependent on the quality, diversity and quantity of training data and the selection of the algorithm. The question of sufficiency of disclosure for inventions involving the application of such algorithms is tied to the scope of the claimed subject matter. A claim directed to an AI achieving a particular effect may be accompanied by a sufficiently detailed characterisation of the training data to provide that effect for the purpose of sufficiency. If a claim purports to solve a problem, the invention must be disclosed in sufficient detail to ensure that implementations of the claimed invention across its entire scope will provide the solution. The emphasis is thus on a careful disclosure of sufficient detail characterising the training and the algorithm in conjunction with a careful claiming of the invention. These are not new challenges for patent applicants across all fields of technology.

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?

We consider that the answer to this question is no for the reasons given above in response to question 10.

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

The question requires clarification. Inventive step is a criterion for granting a patent for an invention codified in S1(b) UKPA and Article 52(1) EPC. The

standard of inventive step is consistently defined in both the UKPA and the EPC as an invention that is "not obvious to a person skilled in the art." Obviousness is assessed with regard to the state of the art[18] and the common general knowledge of the person skilled in the art[19]. The state of the art and the common general knowledge of the notional skilled person with respect to an invention is determined with reference to the date of filing of a patent application for the invention. On this basis, the standard of inventive step is the same between inventions, whereas the state of the art used to assess that standard may be different.

The reference to a "level" of inventive step in the question is therefore confusing, and the question is taken to mean "does AI challenge the *standard* of inventive step required to obtain a patent?".

The answer is no. A new idea discovered with the assistance of AI is suitably assessed for inventive step applying the standard defined in patent law outlined above. Properly applied, this standard recognises how, for an invention in a patent application, any potential for a person skilled in the art to have recourse to AI as a tool is recognised based on both the state of the art and the common general knowledge of the person. It is to be expected that the potential to have recourse to AI may be low in fields of endeavour where the state of the art indicates a low-level or absence of application of AI. It is also to be expected that the potential to have recourse to AI may increase as indications of such are found in the state of the art and/or become part of the common general knowledge of the notional skilled person.

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

The person skilled in the art is characterised by the courts as a skilled but non-inventive person in the art characterised, for example, thus:

> "[They are] deemed to have looked at and read publicly available documents and to know of public uses in the prior art. [They] understand all languages and dialects. [They] never miss the obvious nor stumble on the inventive. [They] have no private idiosyncratic preferences or dislikes. [They] never think laterally. [They] differ from all real people in one or more of these characteristics."[20]

On the face of it, a computer system imbued with the entirety of the state of the art may appear to constitute a reasonable surrogate for the notional skilled person. Such a conclusion results from the legal fiction of the notional skilled person as being aware of the state of the art, not from any innate suitability. One purpose of this legal fiction is to assess what is taken to be obvious, and the characterisation of the person skilled in the art provided above may incidentally align with the capability of an AI. However, it is recognised that the characteristics of the notional skilled person are unrealistic in any natural person, and so constitute a type of *exaggerated* and *unrealistic* person:

---

[18] Section 3 UKPA and Article 56 EPC
[19] Windsurfing v Tabur Marine [1985] R.P.C. 59 CA and EPO Board of Appeal Decisions T 4/98, T 143/94 and T 426/88.
[20] Pfizer Ltd's Patent [2001] FSR 16, at 62

> "A real worker in the field may never look at a piece of prior art—for example [they] may never look at the contents of a particular public library—or [they] may be put off because it is in a language he does not know. But the notional addressee is taken to have done so."[21]

It follows that a machine trained in the art may be useful as an embodiment of the person skilled in the art (for example, for use by patent offices as part of any automation of examination of patent applications). However, an extension to this legal fiction to include *a machine trained in the art* would serve only to expand the fictional capability of the person skilled in the art to have recourse to trained machines (i.e. AI). This expansion is unnecessary. A determination of whether a notional skilled person may or may not have recourse to AI is part of an analysis of inventive step for an invention and is a question to be addressed in the context of the invention, the relevant state of the art and the common general knowledge of the skilled person. It is not necessary to specifically codify any particular tool or technique into the common general knowledge of the skilled person, and doing so renders the legal fiction inappropriate in cases where a skilled person specifically would not have recourse to AI, such as where the state of the art explicitly precludes or *teaches away* from using AI.

Furthermore, the notional person skilled in the art is also relevant for a determination of sufficiency of disclosure of an invention in a patent application. An application for a patent must disclose an invention in a manner which is clear enough and complete enough for the invention to be performed by a person skilled in the art[22]. Thus, to constitute a suitable substitute for, or supplement to, a person skilled in the art, a machine trained in the art would need to be capable of performing an invention on the basis of its disclosure in a patent application. Such an approach is inappropriate and unnecessary. For example, whereas an AI may assist in the discovery of a novel arrangement for a suspension component in a car, a patent for the suspension component may claim the structure and arrangement of the component and how the component is manufactured. For sufficiency it is the component and manufacturing process that must be disclosed in a clear and complete manner. A clear and complete disclosure of such a component and process need not be consumable by a trained AI. It is more likely to be consumable by a 3D-printer or other manufacturing device – yet that need not be specified in the definition of the skilled person because the skilled person would know to have recourse to such a printer to perform the invention. Thus it can be seen that adapting the characterisation of a person skilled in the art to include a trained machine serves no purpose and provides no benefit in respect of sufficiency of disclosure.

Some stakeholders argue that the definition of the skilled person should be adjusted to "the skilled person using AI"[23] including assertions that "it will be difficult for a person alone to come up with anything nonobvious"[24]. Policymakers must not to be misled by such assertions. Any extension of the

---

[21] *Ibid.*
[22] Section 14(3) UKPA and Article 83 EPC
[23] The Reasonable Robot", Ryan Abbott, Cambridge University Press
[24] *Ibid.*

definition of a notional skilled person in this way presupposes that an AI machine is capable of independent invention. AI does not devise inventions. AI is a tool created by, configured by, operating under the instruction of, and with objectives directly or indirectly specified by, a person. Novel ideas discovered with the assistance of AI always involve a human. If an invention is to be determined to be obvious, it must be obvious to a skilled person in the art. The availability (or not) of AI as a tool to such skilled person depends on the state of the art for an invention in a patent application and the common general knowledge of the skilled person.

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

An infringer of a patent is strictly liable under tort law. That a legal person performs an infringing act through the agency of an AI does not affect their liability.

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

Proving infringement of patents is not always straightforward, and in respect of software inventions it can be particularly challenging. Infringement involving the operation of an AI is no different. The challenge can leave some patents seemingly unenforceable for being undiscoverable in use by infringers. The challenge is tempered by high quality professional services engaged in the selection of appropriate inventions for patenting, and in suitable patent claim drafting.

Yours faithfully,

Scott ROBERTS

BRITISH TELECOMMUNICATIONS public limited company