Artificial Intelligence Call For Views: Patents

About D Young & Co LLP

D Young & Co is a leading European intellectual property law firm with offices in London, Munich and Southampton.

Our client base is diverse and international, including academic institutions, blue-chip and global organisations, research institutes, start-ups, and SMEs, spanning all industry sectors and technologies.

We offer a full IP law service, hosting solicitors as well as patent, design and trade mark attorneys within a single legal practice.

Preamble: Definition of AI

The consultation uses the government's definition of AI, as follows:

"Technologies with the ability to perform tasks that would otherwise require human intelligence, such as visual perception, speech recognition, and language translation."

However, we submit that this general statement is ill-suited to patent law, as it defines AI by its technical effect. This creates the risk of either fragmenting the definition of AI according to the field of endeavour, or automatically excluding AIs in some fields irrespective of any other technical merits. Further by defining an AI only by what it can do, it does not attempt to define what an AI is.

We propose that a better definition of an AI is as follows:

"Any technology whose output or functionality is at least in part a consequence of training rather than programming."

This approach captures the unique aspect of AI that it attempts to replicate the natural processes by which intelligence is achieved, in particular that the process involves training or experience rather than programming or hard coded rules. A notable consequence for such AIs is their ability to respond to novel inputs for which no 'rule' has been anticipated; a property referred to as 'generalisation'.

Hence it will also be appreciated that under this definition there is a distinction between machine learning and a program for a computer as such. This is expanded upon below.

A computer program, to the extent that it is a copy-written expression of an underlying idea, is a set of rules authored directly or indirectly (e.g. via a compiler) to achieve an intended outcome; its presence in section 1(2) UKPA originally recognised that circa 1970 computer programs were considered literary works and hence already had protection through copyright. The subsequent evolution of this exclusion through case law is beyond the scope of this discussion.

Some early systems that were called AIs were computer programs of this type, and comprised a sufficiently complex set of rules authored to mimic a notionally intelligent task, whether this was playing chess or capturing an expert's chain of reasoning.

However, a machine-learning based AI system is qualitatively different. As will be appreciated, such a system learns by being trained using examples. Different AI architectures use such examples in different ways, but share the property that the examples contribute to a modification of the AI's functional structure (for example the number and/or strength of neural coefficients). Another property of this training process is that the examples themselves are not preserved within the AI system in a manner that permits independent access; the functional structure of a trained AI system is also not a database.

Rather an AI is different to both; instead of being a rule-based program that may consume data in its operation, or a database that stores data for independent access, an AI learns an output or function by internalising an abstraction of such data within its own structure. It then performs generatively based on this abstraction, either in response to new inputs or spontaneously (depending on the architecture).

As such it is neither an authored work nor a database, neither a collection of rules nor a collection of data, but something new and different to both.

We believe that this consultation is in part a recognition of this distinction, and we believe that the above definition captures it more appropriately for the purposes of addressing the questions below.

Turning now to the questions, each has been answered so as to be read potentially in isolation; hence there is a degree of repetition at certain points.

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

Ultimately the patent system rewards the disclosure of new inventions to the public with a protection of the investment behind that invention, and/or the benefit of exploiting it, for a limited period. This has historically provided the necessary encouragement for further development in other areas of technology and should do so for AI, assuming a level playing field.

Therefore exclusions to this principle need a specific and justifiable reason.

The historic reasons behind the exclusion of computer programs as such from patentability may now be argued to be outdated; e.g. no *sui generis* software right was ever created akin to the database right, and unlike software in the early 1970s, the inventive principles behind modern software are not so clearly bound to a literal expression, which has been relevant to several recent court decisions being unfavourable to program owners. As a result there is no other guaranteed IP right to cede protection to. Meanwhile the software industry is now worth considerably more than many traditional industries and only continues to grow, whilst straddling this apparent gap in intellectual property protection.

It would be desirable for the patent system to avoid this problem again with AI systems, if the opportunity now presents itself.

One can distinguish 'AI technologies' into several fields, including specific applications, architecture, and possibly also the curation of training sets, or the formulation of machine-solvable problems.

The patent system can encourage the application of AI to various fields by acknowledging that in addition to sometimes producing new outcomes or functions, frequently AIs achieve a previously known outcome in a new way that may be simpler, computationally cheaper or just more flexible due to their property of generalisation. Being receptive to the fact that an AI as a technical process has contributed to a known or otherwise apparently pedestrian outcome will help expand the use of AIs in everyday life. See also the final part of our answer to question 9 regarding 'where to draw the line'.

With regards to AI architecture, then from a societal point of view there may be a benefit to protecting even relatively abstract learning architectures, as the ceiling of intelligence for AI is largely dictated by the learning structure. Moreover the existing position that a computer program as such cannot be patentable unless *inter alia* it makes the computer 'a better computer' does not comfortably apply to AI systems, which are themselves a new form of computation. A new AI architecture can itself be a new and better computer for the task that it (and typically only it) performs. Hence whilst it may be easier to award protection for new forms of AI hardware, the patent system would greatly assist AI development by treating AIs, whether in hardware or software forms, as emulations of learning systems rather than as programs as such.

The patent system can also provide encouragement by protecting one of the major sources of cost and investment in AI, namely the creation of training sets. Whilst in principle specific training sets may be protected by copyright or a database right, the generative principles behind certain useful training sets may be worthy of protection. Whilst there is not an existing exclusion in statute that acts as a barrier to protection, a positive approach to such inventions would enable safe investment in large scale training sets, either by creating new data or curating / filtering existing data.

In short, with the appropriate will, there is scope for the patent system to treat AIs differently and more generously than, in particular, computer programs as such, and recognise them as a distinct form of computation, either when disclosing new computational approaches or when applied to new or old problems.

Question 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?

As per the consultation description, we assume that the AI at issue is not considered conscious or a person in any sense. Hence we assume the consultation does not relate to so-called 'strong' AI.

2a: Can current AI systems devise inventions? Particularly, to what extent is AI a tool for human inventors to use?

To the extent that an AI has no independent agency, then it must be considered a tool as a matter of definition.

This is not to say that an AI is not capable of generating things that are not obvious to a person skilled in the art, as per the requirements of section 3 UKPA. However, regarding the specific wording of the question (and section 7(3) UKPA), it is not clear that an AI is a <u>deviser</u> of an invention. There is an implication of planning and forethought within the term 'devise' that is absent from current AIs.

Typically, the planning and forethought are provided by the parties who curate a training set for domain specific AIs, use a specific AI architecture, and/or construct/convert a problem into a solvable form by an open-domain AI (see our answer to questions 2c and 6 below for a further discussion of this). The AI then generates an output or function that may be new and non-obvious.

This generally points toward the conclusion that an AI is a tool.

Please also see the answer to question 3 below, 'Second Approach'.

2b: Can current AI systems devise inventions? Particularly, could the AI developer, the user of the AI, or the person who constructs the datasets on which AI is trained, claim inventorship?

Please first see our response to question 3. In short, we would suggest that non-conscious AIs are capable of discovering new and non-obvious things already latent within the domain of their training set and inputs, and the industrial application of these discoveries by legal persons are in principle inventive.

With this in mind, we address the above question.

If an AI could be identified as a sole inventor, then it follows that there may be cases where a human cannot claim inventorship.

However as noted in our response to questions 2a and 3 and as summarised above, it may be preferable to consider an AI to be a tool for discovery.

In this case, in an analogous fashion to other, sometimes automated, discoveries (e.g. for genes and drugs), an industrial application of the AI's discovery may be considered inventive. In this case, as appropriate any one of the developer, user, or curator of the training set may claim an inventive industrial application of the AI's output or functionality.

As is also noted in relation to question 6 below, it will be appreciated that the person who has made the arrangements necessary for the production of the invention may not necessarily be the person who built or trained the ML.

A person who created and/or trained an ML is likely to be the inventor in the case of domain-specific MLs, whose function is defined primarily by a curation of the type of training materials, if not the specific training set.

By contrast for an open-domain or generalised ML the inventor / owner is likely to be the person who, in appreciation of a problem that needs solving, constructs a task for the open domain ML to solve.

In other cases, the architecture may be critical to the desired output or function.

Hence inventorship will be a function of contribution as much as it is now for any other field of technology; but critically, the field of technology itself should not exclude the people who made the arrangements necessary for the invention to arise from claiming inventorship.

2c: Can current AI systems devise inventions? Particularly, are there situations when a human inventor cannot be identified?

If an AI produces new outputs or functionality that cannot be devised or realised by conventional means (e.g. by programming), then there is an argument that there may be circumstances where the developer, user, and/or trainer of the AI cannot be said to have devised the specific output or function.

However, it is also clear that they did make the arrangements necessary in order for the AI to generate the specific output or function. As such they did devise the general nature or domain of the output or function and typically its relevance to the problem at hand, but allowed the AI to express a specific solution in response to the arrangements that had been made.

Put another way, if an AI is trained multiple times, it may express different specific outputs or functionality, but within bounds anticipated and devised by the developer / user / trainer, and typically addressing the same problem in a similar manner.

As noted in our response to question 6, an AI that is trained to distinguish cats from dogs will never generate a new protein model or circuit layout; the creator of the AI has determination over the domain of the AI's application (or equivalently the user has determination over the behaviour of an open-domain AI).

Consequently it can be argued that there is never a case where the output or functionality of an AI is not ultimately governed by a person, or falls outside of the bounds of their reasonable expectation; where this output or functionality solves a particular problem, it may be reasonable to argue that the person has invented the solution, using the AI as the means.

Nevertheless, as noted initially, this does not preclude the possibility of an AI generating something without human involvement that plausibly meets the criteria for inventorship – but based on the above arguments, this should be the exception to be proved rather than the default assumption.

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

As per the consultation description, we assume that the AI at issue is not considered conscious or a person in any sense.

There are at least two approaches to this question, which in part depend on the answers to other questions within this consultation.

i. First Approach

The first approach illustrates some potential problems with considering an AI as an inventor.

For the first approach, a preliminary matter is to define what an invention is. For the purposes of patentability in section 3 UKPA an invention is such if it is not obvious to a person skilled in the art, having regard to any matter which forms part of the state of the art.

Hence an invention is defined on its own merits and without regard to the status of the inventor. Matters of ownership and the right to grant of a patent, which require a legal person, are separate.

There is no doubt that AIs can produce things that are new and not obvious to a person skilled in the art. Hence it would appear that there is no legal barrier to an AI being an inventor based on a consideration of the invention alone, and on this basis the answer to the question appears to be yes.

If this position is taken then we note, with reference also to question 6, that there is a clear analogy with section 9(3) of the Copyright, Designs and Patents Act 1988 (CDPA), where if a machine creates a work then authorship is accorded to the person 'by whom the arrangements necessary for the creation of the work are undertaken'. In this case, a corresponding rule could accord inventorship to the person who has made the arrangements necessary for the production of the invention (this person is discussed in more detail in the answer to question 6).

However, moving on from this initial position, a further matter is to ask whether there is any other, possibly implicit, requirement of an 'inventor'. If one additionally considers the act of invention to require insight and the recognition by the inventor that they have conceived of an invention (at least to the extent that a patent is being sought), then for the AIs at issue the answer may be more nuanced.

As noted above an AI can generate things that appear new and inventive, but conventionally they could not recognise this fact; there is no selectivity. An AI could in principle however generate an estimate of the potential inventiveness of the generated thing, based for example on measurements of likelihood or difference from internal values or comparisons with prior outputs. If this is considered an adequate facsimile for insight and recognition then the answer may arguably be yes. Meanwhile if not, or if it is not implemented in a particular AI, then the answer may be no.

The likely outcome of this approach is a fragmentation of the qualification for inventorship that depends upon aspects of ML operation that are independent of the actual invention itself and which, due to the differing nature of some MLs suited to some tasks, may result in arbitrarily different inventorship outcomes for different fields of endeavour. This is clearly unsatisfactory.

Pushing this issue further, the requirement of an inventor in law can be derived from section 7(3) UKPA, which requires that an invention is 'devised' in order to be accorded to an inventor; e.g. in some manner the invention is intentionally sought out. If this is the bar set for inventorship then it appears unlikely that current AIs meet the criterion. This appears to be the position taken by the UK High Court in the recent 'DABUS' case (*Thaler v Comptroller* [2020] EWHC 2412 (Pat)).

However, levelling the playing field by stating as a matter of principle that a non-conscious AI cannot devise an invention, so as to side-step this problem without resolving it, also appears unsatisfactory and will stifle innovation if it blocks a route to patentable rights for such inventions.

Hence in summary this first approach can give rise to different answers depending on different initial assumptions. Firstly, if inventions are only assessed on their merit as per section 1(3) then the answer may be yes. If however it is considered that an inventor must also have an air of invention about them to qualify, then the answer may be yes or no depending on how and to what extent the ML implements identification or recognition of an invention, separate to the invention itself, and this can have unjust consequences. Meanwhile if an inventor must have the agency required to intentionally devise an invention, then this appears to rule out existing AIs by fiat.

It will be appreciated however that even if this first approach is adopted, then regardless of whether the AI is an 'allowed' inventor, the patent system can still allow the granting of patents if it is accepted that the person(s) who arranged the circumstances under which the invention arose may also be eligible to be named the inventor or owner (see also questions 2a and 6).

Nevertheless, a better approach may be preferable.

ii. Second Approach

As an alternative second approach, we suggest that an AI may instead be treated as a means to discover new and non-obvious properties that are latent within a space occupied by the AI's training set and inputs. In this sense, an AI may discover, but not invent, a new thing.

An AI creates an internal representation of features of its training set, and so the information available to the trained AI is thus typically a partial and transformed representation of the training set, as determined by the type of AI. Sometimes an input may also be used by the AI as a baseline or scaffold for using its internal information, and hence a stimulus for the AI's output or functionality. Hence all the AI's possible outputs are a function of the internal information derived from the training set and optionally the current input. These have been provided, arranged, or caused by the curator of the training set, the AI architect, and/or the user, and are latent within the trained system. The outputs or functions of the AI are thus explorations of this latent space and can be better thought of as discoveries.

In this case, in a manner analogous to existing case law relating to drug discovery or gene discovery, an industrial application of the discovery made with the AI may then be inventive. The bar for this industrial application can be very low since by definition being based upon the discovery it will be new and non-obvious, if the discovery is. Hence a claim making such use of a new and non-obvious output or functionality of an AI system should be an invention (leaving patentability as a separate issue).

Hence in this case the AI is considered a tool for discovery, albeit a seemingly creative one, and the industrial application of its discovery (e.g. by a person) is inventive.

This approach avoids most of the pitfalls relating to issues of inventiveness raised in this consultation.

In short, focussing on the industrial application of generated output or functionality of an AI, which is treated as a discovery, is a more stable and consistent approach than focussing on the input, training, architecture, intentionality or framing of an AI as a means of divining inventorship, both now and as AIs become more capable and potentially less straightforward to analyse.

Firstly, this approach is independent of the type of ML being used. This avoids any issues relating to *how* an ML invented or identified an invention. Notably it also avoids the need for an 'explainable' AI that reveals the specific path of determination in producing the result, any more than the path to a conventional discovery is necessarily explainable before it is industrially applied.

Secondly it is consistent with case law on existing automated discoveries, and thirdly it avoids the need for any new legislation or rules relating to AI inventorship or co-inventorship, or for any parallel issues such as who the corresponding skilled person in the art may be.

In conclusion therefore we would suggest that non-conscious AIs are capable of discovering new and non-obvious things latent within the domain of their training set and inputs, and the industrial application of these discoveries by legal persons are in principle inventive. We believe that this is the best approach to the determination of what has been invented and hence also who is the inventor.

Question 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?

Addressing these in turn:

4a. If AI cannot be credited as inventor, will this discourage future inventions being protected by patents?

Only if no-one else can be credited as the inventor instead. As per our replies to questions 2, 3 and 6, we believe that the person who made the arrangements for the AI to generate a novel and non-obvious output or process can typically be called the inventor, at least of an industrial application of the output or functionality of the AI.

4b. Would this impact on innovation developed using AI?

We note that the question says 'developed using AI' rather than 'developed by AI', and in accordance with our other replies in this consultation we think this is the key point: the AI, given a lack of agency, is a tool for innovation rather than an innovator itself.

Therefore, we believe that by simultaneously treating AIs only as tools for discovery and also allowing inventorship to be conferred on those who industrially apply those discoveries, innovation can continue to flourish.

Meanwhile if the law does not allow a route to patent ownership from the generation of new and non-obvious things by AI, then this creates a clear disincentive to use AIs in research or to apply for patent protection.

4c. Would there be an impact if inventions were kept confidential rather than made public through the patent system?

Clearly yes; trade secrets frequently result in the unnecessary repetition of work amongst competitors, and are often inherently insecure as the basis for significant investment in innovation either because they leak or because the solution can be independently rediscovered later. Furthermore there would be a loss of transparency about how potentially important AIs within society operate. The entire point of a monopoly patent right is to address the impact of trade secrets.

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

There is no moral case for rewarding an AI itself, in the sense that it would not be injured by the lack of recognition. Meanwhile by contrast rewarding an AI may deny recognition for any contribution of those

who created the AI, or recognised the applicability of what it generated to a particular problem. Hence there would appear to be a bias against recognising an AI on moral grounds.

Separately, we note that morality or ethics in AI is a broader issue than one of inventorship.

One commonly quoted issue is that of AIs in autonomous vehicles being faced with ethical dilemmas (e.g. if an elderly person walks in front of a vehicle, should it swerve one way and likely kill a motorcyclist, swerve another and possibly kill the driver, or stay on course and certainly kill the jaywalker? i.e. the so-called 'trolley dilemma'). Such dilemmas are not limited to cars, but are also being considered for care robots and other AIs such as medical, social (e.g. insurance, loans etc.), and judicial screening systems.

An emerging consensus with these systems is that there should ultimately be someone accountable for the AI's actions; a company cannot wash its hands of its device's actions on the basis that they were performed by an AI. In short, a legal person will be responsible for an AI's behaviour (see also question 15 below).

In this sense it appears that there will be an assumption in law that an AI's behaviour will be considered to be, if not transparently or auditably deterministic, then at least within a boundary of expectation for which the person responsible for the AI is culpable.

This would appear to contradict the position that an AI can also produce ideas that are *outside* the boundary of expectation for which the person responsible for the AI is 'culpable'; i.e. an invention by the AI for which that person could not be considered to be at least a co-inventor.

In short, it appears difficult to coherently accommodate in law both accountability for, and inventiveness by, an AI.

Hence in anticipation of a broader legislative stance on morality and ethics for AIs, it would appear prudent to assume by default that the output of an AI is at least within a boundary of expectation of the person(s) responsible for the AI, and that as a result there should be at least a *prima facie* assumption that such persons are at the very least co-inventors.

Similarly by analogy with section 43(3) UKPA, to the extent that an AI may be considered to 'assist' with an invention framed by the people responsible for it, the AI should not be accorded the status of an inventor.

It will be appreciated that our position as outlined in other answers herein is that the AI is not an inventor at all but a discovery tool, and that in fact the invention lies in an industrial application by a legal person of any discovery made by an AI. This would also be consistent with the above anticipated stance on morality in AI.

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?

Our position in questions 2 and 3 is that a human should always be able to be named as inventor, as the AI is merely a tool for discovery. Consequently ownership operates as per current statute.

However, if an AI *was* identified as a sole or joint inventor, then as noted previously, s9(3) CDPA may be a suitable model to determining ownership. The analogy with s9(3) CDPA) for inventorship (first mentioned in relation to question 3 above) also be applied to ownership, since in the first instance the author in copyright is also the owner.

Hence if an AI is recognised as the 'creator' of an invention, then a first option regarding ownership (whether or not the AI is named as sole or joint inventor) is to use s9(3) CDPA as a model. This provision covers authorship of works by computers (e.g. weather maps, which can be created with no human intervention but have commercial worth): "In the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken".

In this case the arrangements necessary for the creation of the AI may for example be the training set and any process used to create the AI, which will have been selected to determine what the AI does; it will be appreciated for example that an AI trained to distinguish cats from dogs will never generate a new protein model or circuit layout; the creator(s) of the AI have determination over the domain of application by the AI. Hence a person will have made the arrangements necessary for the production of the invention even if they did not anticipate exactly what it was, as per the computer generated works in s9(3) CDPA. This would seem to be the obvious way to resolve the issue.

It will be appreciated that the person who has made the arrangements necessary for the production of the invention may not necessarily be the person who built or trained the AI. The specific contributors to a given invention will be a question of fact.

Hence as noted above a person who created and/or trained an AI is likely to be the inventor / owner in the case of domain-specific AIs, whose function is defined primarily by a curation of the type of training materials, if not the specific training set.

By contrast for an open-domain or generalised AI the inventor / owner is likely to be the person who, in appreciation of a problem that needs solving, constructs a task for the open domain AI to solve in a desired manner. An example of an open domain AI is GPT-3, a system that has been trained on any and all text within large public databases to act, in effect, as a generalised predictive text engine; the breadth of topics it has trained on means that it can produce results that address newly posed problems (see for example https://www.theverge.com/21346343/gpt-3-explainer-openai-examples-errors-agi-potential).

Meanwhile even for self-training systems such as AlphaGo Zero, the ability to experiment and observe for the purpose of self-training, and the domain in which that experimentation occurs, is provided and enabled by the developers of the system.

Hence we believe that if necessary a new rule similar to the provision in s9(3) CDPA could address the issue of entitlement to own a resulting patent, even if an AI was considered a sole or joint inventor.

Such an approach could avoid a sequence of issues that could arise from the AI being treated as an inventor but then falling foul of subsequent laws not just in the UK but also overseas. For example under US law, a route to ownership by a separate legal person could enable prosecution to continue in the clearly likely event that the AI inventor is incapable of providing an assignment and declaration.

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

There appear to be two main themes within the patenting of AI, relating to whether a system comprising an AI is a means to an end (in which case the issue is mainly inventorship/ownership of the resaults) or an end in itself (in which case the issue is mainly patentability of the AI).

The first theme is dealt with in questions 3 and 6 above. Our position is that AIs are not inventors but their discoveries can be industrially applied as inventions by legal persons. However even if AIs are deemed inventors, there appears to be an apt analogy with the route to ownership of AI creations found in copyright, and/or a likely need to identify a co-inventor for consistency with laws on accountability for AI behaviour in other fields.

Meanwhile the second theme relates primarily to excluded subject matter, and is dealt with in more detail in question 9 below. Our overall position here is that AIs frequently do not fit the categories they are pressed into during prosecution, and that in particular an evaluation of AIs as a separate category to computer programs (e.g. an emulation) would be of considerable benefit. Any approach that excludes AIs from patentability (whether in terms of training, architecture or application) or makes their protection uncertain would be an unnecessary and destructive problem for UK industry, particularly if it was more proscriptive than protections given elsewhere in the world.

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

Again as per the consultation description, we assume that the AI at issue is not considered conscious or a person in any sense.

Clearly a qualitative change in AI (e.g. to so called 'strong AI') could necessitate a further consideration of patentability.

For now though we do not believe that incremental improvements in AI will affect patentability if the legal principles are well formulated from the outset.

Question 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?

We address the excluded categories in turn below:

(a) a discovery, scientific theory, or mathematical method, as such:

An AI system per se is clearly not a discovery or scientific theory as such, but in principle could produce either of these; if so the result can be treated on its merits like any other discovery or scientific theory, for example benefiting from a claim to an industrial application. As will be appreciated, this is the theme of several of our other answers.

An AI system *per se* is also not a mathematical method is such. A mathematical method is a defined process for performing a calculation. By contrast, and as per the definition provided above, rather than

being a rule-based process, an AI learns a function by internalising an abstraction of training data within its own structure. It then performs its function generatively based on this abstraction, whether in response to new input or spontaneously (depending on the architecture).

As such it is inherently a thing, and its functionality is embodied within it. No mathematical method can be detached from the AI and asserted to be protected in an abstract sense. Consequently any claim comprising an AI or the operation of an AI is inherently not a mathematical method as such.

By similar reasoning, a claim comprising an AI or its operation is also not a scientific theory.

(b) a literary, dramatic, musical or artistic work or any other aesthetic creation whatsoever, as such:

Beyond a literal software copy of it, an AI system is not an LDMA work as such. However AI systems clearly can create LDMA works. For such works, copyright law and the exclusions in section 1(2) UKPA can apply as appropriate.

(c) a scheme, rule or method for performing a mental act, playing a game or doing business, or a program for a computer, as such:

We address each of these in turn.

- *i. Mental act as such*

There is a clear temptation to suggest that the technical effect of an AI is the performance of a mental act; this temptation may be reinforced by the government's effect-focussed definition of AI, which as noted in our preamble above appears ill-suited to a discussion of patentability.

We note that the purpose of the mental act exclusion is to ensure as a matter of policy that people are free to conduct their thoughts, in a similar way that the purpose of the treatment and diagnosis exclusions of section 4A UKPA are to ensure as a matter of policy that people are free to conduct medicine.

But in the same way that section 4A UKPA does not prevent the patenting of tools that may be essential for such medical practices, section 1(2)(c) UKPA should not prevent the patenting of tools that produce similar effects to mental acts.

Specifically, we submit that a claim to an AI that performs text recognition, for example, cannot be a claim to a mental act, and does not prevent or encompass the mental act of text recognition, precisely because it is a claim first and foremost to an AI. It is only by using an AI that performs text recognition that a third party could infringe the claim; therefore the effect of the claim is not a mental act as such and does not limit how people conduct their thoughts.

It will also be appreciated that typically an AI performs this functionality in a way that is fundamentally different to the performance of a mental act by a person.

Hence we would argue that a claim to an AI should not be excluded simply because the same overall goal or task can be achieved by a mental act; the mere fact that a claim utilises an AI should be sufficient to ensure this. Any further features that make clear that the AI operates in a different manner to a

conscious thought process then simply provides further clear distinction – and, frequently, a different actual result when clarified.

- *ii. Playing a game as such*

To the extent that an AI is involved in the operation of a scheme, rule or method for playing a game, it is no different to any other contributory component.

We do note however that where AIs are used to accommodate or provide a rich environment within a game, such as handling new combinations of unscripted circumstances (where the permutations of such circumstances are too many to hard-code), the AI is providing the ability to have these combinations occur, over and above any notion of a scheme, rule or method of playing a particular game as such. It will be appreciated that the exclusion originally related to rules that define the play and outcome of a particular game, such as the allowable moves in chess and the conditions for winning. However, ensuring that a non-player character points in the right direction or selects a meaningful response to the user in a new and unanticipated combination of circumstances, or the like, should not fall under this exclusion where they relate to maintaining the fabric of the game, and not the rules by which it is played. Similarly AI systems that assist a user may operate within the specific rules of the game, but do not necessarily define them. The applicability of this exclusion must therefore be considered with great care.

- *iii. Doing business as such*

To the extent that an AI is involved in the operation of a scheme, rule or method for doing business, it is no different to any other contributory component.

- *iv. Computer program as such*

It will be appreciated that the definition of this exclusion is one of the most fluid in patent law, but we are working with the current definition in case law as per the five signposts in *Symbian*, whilst recognising that these are merely guidelines and not a definitive test (MOPP 1.37.1-2).

As noted in our proposed definition of AI, there is a distinction between a rule-based computer program and an AI that has learned a function by internalising an abstraction of its training data within its own structure.

As a result, we believe that it should be assumed *a priori* that an AI is evaluated in a different manner to a computer program (even if some considerations may sometimes converge).

An AI as defined previously is a new and unique approach that fuses computation and external data within its own architecture; the result is typically a single-purpose computing system with a beneficial ability to generalise its behaviour in response to new inputs.

Hence one property in particular that distinguishes an AI from a program as such is that, when trained, it embodies external data (for example, real-world training data) within its structure to create a new and typically bespoke form of computer. This may be seen as consistent with patentability from the decision in *Macrossan*, in which the court rejected the application on the grounds that system had not contributed a new form of hardware, unlike in the parallel *Aerotel* case.

Hence where an AI may be considered to be a new form of computer, it may be considered patentable even if it is only emulated in software.

In this sense an AI is perhaps more akin to a field programmable gate array, where functionality is encoded physically in the architecture of the system. Indeed it will be appreciated that there already dedicated neural chips for a similar purpose. In this sense, a software implementation of an AI is an emulation, not a program. Indeed, it will also be appreciated that most trainable AIs use artificial neural networks, i.e. they are an emulation of a (typically hypothetical) arrangement of neurons.

Meanwhile *Symbian* asks whether a computer program reveals a technical contribution, providing five possible and non-exhaustive indicators of this.

Clearly an AI may be capable of demonstrating any of the five signposts, depending on the specifics of the AI.

However we would suggest that, in many circumstances, it should be recognised that it is not appropriate to apply the signposts to an AI claim at all, on the basis that it is an emulated bespoke computer, rather than a program for a computer. Consequently the signposts can fail to address the features of the AI that are of value.

In particular, where the effect of an AI is itself seemingly non-technical (for example reading, navigating, drawing or the like), the fact that a non-human system has been created to implement the task at all is often the desired and valuable technical effect. The current signposts are not set up for this.

Even if the signposts are applied, this should be done with caution.

It will be appreciated that the operating portion of an AI is not a program; it has not been authored in the conventional sense and does not operate by conventional program rules. As a result it provides technical effects within the software environment that it operates in that may be interpreted as making a computer that operates in a new way, or making a better computer, or being an emulated co-processor of a computer that provides newly flexible responses to messy inputs. It is not merely a specific sequence of operations of the host computer itself but an emulation of something else. Hence even if an AI is not treated as an emulated bespoke computer or processor in its own right, there may be a presumption that an AI improves the operation of a host computer.

It will also be appreciated that a trained AI system embodies external (typically real-world) information. As such the AI can be the direct product of a technical interaction with real-world features. In a similar manner to how inputs and interactions with computers are frequently patentable, it may be considered that such an AI can incorporate 'pre-experienced' interactions and external stimuli to a device. This may confer similar patentability rather like a reverse *Vicom*, where an external effect has been used to modify the behaviour of the computer, by embodying the effect within it.

In short, we believe that because AIs are not programs for computers *per se*, there are more routes to patentability for AIs than under the current regime for computer programs as such, and this should be reflected when evaluating patent application for such systems.

(d) the presentation of information, as such:

A machine learning system *per se* is clearly not a presentation of information as such. Furthermore, as noted above even its internal representation of training data is not in an independently recoverable form.

However of course in principle an AI may cause a presentation of information, and to the extent that this falls under the exclusion as such it may be no different to any other contributory component of a system.

However again we would caution that where an ML is able for example to provide new information, or provide information in more challenging circumstances, then this should be distinguished from the presentation of information as such, as appropriate.

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

The discussion of excluded subject matter above relates to the boundaries of technical effect; it will be appreciated however that one area where AIs may have their greatest use is in replicating human activities that are seen as *prima facie* non-technical. Whether the activity is reading, speaking, driving a car, walking, sorting, or any other field, we believe that the presence of an AI within a claim should be a signal to additionally evaluate the fact of implementing the effect, in addition to any technical aspect of the effect itself. Put another way, teaching an AI robot to play a piano could be dismissed as non-technical because the result is the playing of a piano; however, the line should be drawn earlier – one technical effect is actually producing an AI that plays the piano in the way that *it* does, and not necessarily the playing of the piano itself.

We have also noted the potential societal benefit of protecting even relatively abstract learning architectures, as the ceiling of intelligence for AI is largely dictated by the learning structure it uses. This may require a reappraisal of exclusions such as mathematical methods, computer programs as such, and the like, in light of a claim to an AI having a particular makeup. Again, treating AIs as emulations may be beneficial in this regard.

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

As with answer 4c above, if an AI is not patented it may instead be kept a trade secret; this will certainly be in conflict with any notion of ethical oversight.

Even if oversight may then be conducted *in camera* to preserve the trade secret, this still prevents transparent ethical oversight. Furthermore the entire system is fragile as any leak of a trade secret may destroy its value, and this could stymie cooperation between AI developers and any oversight body.

Furthermore, such ethical oversight is likely to be the remit of parties unrelated to patents, and who may not appreciate such issues or may weigh other requirements of transparency more heavily. This may discourage investment in AIs if mandatory disclosure is possible and patent protection unobtainable.

Given the increasing use of AIs in medical, legal, administrative, transport, and industrial settings, a right that reliably gives AI developers the confidence to disclose and discuss their systems seems essential.

This is particularly important for AI since a specific AI cannot necessarily be reverse engineered or decompiled for independent inspection.

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?

This question includes the following notes:

In answering this question, you may wish to consider:

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

These are addressed as follows:

Sufficient disclosure is a requirement and a problem for any invention, and generally the principle is that 'enough detail' will vary with the complexity of the task and/or the predictability of the asserted outcome.

It is a fact of life however that sometimes this detail is not fully known in advance, which is why the courts are equipped to subsequently demand evidence that the description does or does not properly cover the scope of the claim.

For example, a patent could be granted for an application disclosing a catalyst that works with compounds containing X, within a certain range of concentrations. It may then transpire years later that for compounds containing X *and* Y, the catalyst no longer works in this range. As it stands the claim therefore comprises non-working embodiments.

However, quite rightly there is no current requirement that the patent should therefore have demonstrated that *every* possible compound comprising X works with the catalyst in order to grant the patent in the first place; merely that this is a plausible outcome.

Similarly, a patent that says a shallow and fully connected neural network is trained to distinguish dogs and cats can be plausibly expected to do this. Indeed in this case the skilled person may also appreciate that such a system may make some errors and still be within the scope of the invention, as this is

inherent in the nature of such AIs. Hence this is enough detail for these circumstances. There should be no requirement to disclose every possible / appropriate training and testing image.

For more complex AIs, or more complex tasks, the additional complexity of the system typically only exists in order to address a specific need, and so the relevance and contribution of these additional complexities to the invention will be explained in the same way that additional complexities of any other invention are disclosed; the burden of understanding imposed on the skilled person is the same.

The yardstick is thus whether the explanation contains enough detail to plausibly result in an AI that falls within the scope of the claim, and the amount of explanation provided will vary depending on how many features/conditions are critical in achieving that outcome. This is ultimately the same requirement as for any other invention.

If it is subsequently shown that there is an outcome that conflicts with the claim scope (e.g. a failure when using training sets X *and* Y), then this can be handled in the same way as any other patent, by being revisited in the courts or with a validity opinion by the UKIPO.

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?

Referring back to questions 4c and 10, we believe that patents should provide AI developers with the confidence and security to subsequently disclose their inventions in a manner that benefits society; however to the extent that this disclosure goes beyond the current purposes of patent law, this would be a matter for public scrutiny or separate regulation relating to the specific issues raised.

Meanwhile, we do not believe that such considerations should affect the requirements for disclosure for the purpose of obtaining patents themselves, for a number of reasons.

Firstly, in the case of already well documented societal issues such as inadvertent racial or gender bias in AI systems that have been acquired from training data, there is no obvious way in which this could be identified from a written description of the AI system per se, or necessarily even the training data; indeed, if it could then the developers would have likely rectified it already.

It will also be appreciated that issues of unknown or indeed known bias exist in other fields but are not considered an issue for inventions in those fields; for example the bias in crash test data: https://www.consumerreports.org/car-safety/crash-test-bias-how-male-focused-testing-puts-female-drivers-at-risk/.

Furthermore, these issues generally do not read on to novelty or inventive step, and so their relevance to a patent Examiner appear limited.

Even taking account of sections 1(3) and 1(4) UKPA (for an invention that contravenes morality or public policy), it will be appreciated that these are only relevant during examination to the extent that the problem is *prima facie* obvious to the Examiner; subsequently arising instances of the problem are likely better handled by the courts or independent regulation.

Secondly, there are severe practical issues relating to sufficient disclosure. Most AI systems will be characterised by their training data, their architecture (and the resulting coefficients / internal trained structure) and the test inputs, generated outputs, or the manner of framing these.

In a 'big data' application, the size of the training set may be enormous. For the GPT-3 example mentioned elsewhere herein, the entirety of Wikipedia (6 million articles) represents just 0.6% of its training data. It would be impossible to disclose such a training set in a patent application (even if provision was made for an accompanying data file similar to a gene sequence file for a biotech patent). Furthermore, doing so would be unlikely to achieve the goal of the question, since it is only through operation of the AI itself upon this training set that some of the biases within it become apparent.

Similarly, a table of the raw coefficients or weights of an AI would be of little use in assisting the Examiner with assessing whether the disclosure is sufficient. The Examiner is not going to be able to tell from a table of raw coefficients or weights, or a set of training data, whether the model works. Furthermore, it only discloses one instance of the system, whereas the claim may be directed to a broader principle. Hence for example a disclosed training set of Siamese cats and bulldogs does not guarantee that a claimed network trained on 'cats and dogs' can distinguish them to a sufficient degree. Even if a fully trained system comprising a complete training set, training order, coefficient table, test set, and test results, was provided as a disclosure, this would still only prove that one instance of the AI performed as described.

In this regard it will be appreciated the coefficients/weights themselves are typically not the core of the invention. If the same training examples are processed in a different order, this will likely lead to different coefficients/weights, but with both defining working models within the claim scope (e.g. cats v dogs). The invention lies in the principles involved in arranging for the model to carry out a particular task, not in a specific instance of such a trained model.

Hence, and as with question 11, a disclosure should be sufficient if it is plausible that certain disclosed features or properties of the AI system support and enable the claim, with the rigour required for plausibility relating to the complexity of the situation and the scope of the claim.

Thirdly, even for those AI models where disclosure of the training data or coefficients/weights is practical, the training data and coefficients/weights may be considered commercial secrets that have taken considerable investment to develop, so an applicant is unlikely to want to disclose this. If such disclosure was required, the consequence could be that applicants no longer seek patent protection for AI inventions, effectively denying patent protection in a field of technology contrary to the TRIPS requirement that patents should be available in all fields of technology. Furthermore, the requirement of sufficiency of disclosure merely requires that sufficient information is made available that a skilled person in that field can implement the invention. It is not a requirement that information is made available so that another person can implement the invention without *any* development effort. There is no need to give a competitor a short cut to avoid incurring the development costs for implementing an inventive concept that is otherwise sufficiently described.

Fourthly, details relating to a training set, weights or coefficients, or a set of sample outputs, would impose a burden upon AI patents that is not imposed on other inventions; namely that a working prototype of the invention must be constructed prior to filing the application. After all, it would not be possible to disclose a training set, weights or coefficients, or a set of sample outputs, without

implementing the AI first. Notably it is not a requirement in other fields to provide such a working implementation of an inventive principle, such as software source code for a software-implemented invention, or an integrated circuit layout for inventions relating to hardware inventions. Typically again the embodiments are described to an extent that plausibly enables the invention within the claim scope.

For at least these reasons, we believe that the disclosure of a patent should be limited to a disclosure that plausibly enables the skilled person to reproduce a system that falls within the scope of the claims, after suitable developmental effort, as is the case for other fields of invention.

Any further disclosure is for a separate purpose and should be conducted in a separate forum.

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?

We note that in our other answers, we have suggested that an AI should not be considered an inventor. As such we would not seek to answer this question from the perspective of the degree of inventiveness of an AI, or the extent to which an AI contributes to an invention. We do not believe that this is a fruitful line of inquiry. However for more on this topic, please also see related question 14.

Rather, we suggest that AI systems are new tools for achieving new things, and that the pursuit, recognition, and industrial exploitation of these new things can be treated as inventive in just the same way that inventions based on earlier enabling tools are assessed on their merits.

When a new technique appears (for example, CRISPR enabling gene editing), then new applications come within reach that were not practical before; in this sense the hurdle of gene editing *per se* has been overcome by the previous inventive step of CRISPR itself, and so the novelty or inventiveness of a new application within the field of gene editing will be evaluated after taking this into account.

In a similar way, where a new technique appears (such as deep learning within AI), then new applications come within reach that were not practical before; but using deep learning as a baseline, the novelty or inventiveness of the application itself should then be assessed on its own merits.

In short, known AI techniques contribute to the state of the art, and new applications or techniques can be evaluated against them, like any other field of invention.

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

We note that in our other answers, we have suggested that an AI should not be considered an inventor. In the sense that an inventor should be evaluated against their peers, if this position is adopted then it would appear reasonable to conclude that the person skilled in the art should not be extended to the machine trained in the art, as a matter of consistency and principle.

We also note that the existing ‘person skilled in the art’ recollects the art, either singly or in combination. However, as noted in our original definition, an AI is not a database for recollection but instead its outputs are inherently generative and generalise; a machine trained in the art would not recall the art but would generate something else based upon it. This would inevitably extend the notion of novelty and furthermore make it inherently unclear, as the notional machine trained in the art does

not exist and so what it might produce is speculative. Meanwhile if, to avoid this, it is decided that the machine just recollects the art, then there is no point including it over and above the skilled person.

Furthermore we believe that we should be slow to deviate from existing patent law to accommodate any new technical development, including machine learning; the majority (if not all) of the principles of patent law continue to apply.

In particular, case law recognises that the skilled addressee may use tools, and it follows that it is likely to recognise a machine learning algorithm as one such tool in due course (if not, there may be the need for legislation). The concept of the "machine trained in the art" is too fluid, with the definition, capacity and standards attributed to a "machine" subject to change as AIs change and improve.

By contrast, a "person" has inherent and well established limitations and can be treated as a fixed standard.

Referring to Jacob LJ's description of the skilled addressee in Technip France SA's Patent [2004] EWCA Civ 381: "It is settled that this man, if real, would be very boring – a nerd. Lord Reid put it this way in [Technograph]: '… the hypothetical addressee is a skilled technician who is well acquainted with workshop technique and who has carefully read the relevant literature. He is supposed to have an unlimited capacity to assimilate the contents of, it may be, scores of specifications but to be incapable of a scintilla of invention… The man can, in appropriate cases, be a team – an assembly of nerds of different basic skills, all unimaginative ...'

An AI system, as a tool, may help to make more real the legal fiction of the skilled addressee set out above (who is deemed to have an unlimited capacity to assimilate information), in much the same way that the internet has done; but the skilled addressee remains unchanged as new tools become available.

It will of course be appreciated that the skilled person may become *more* skilled as a consequence of having better tools, and as the available art grows; for example the skilled person as at 1900 in the field of aviation is not the same skilled person today. However this is still a property of the skilled person as evolved over time, and does not require any new and independent yardstick for inventiveness based on trained AIs.

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

As set out above, a tool (e.g. an AI system or algorithm) does not and should not attract legal personality. Rather, only individuals and corporations should do so. In that vein, an AI system should not be capable of infringing a third party patent. Rather, it should be a person or corporation which 'makes arrangements' to implement said AI which should be the infringer.

In any event, in circumstances where a patent by definition must firstly have industrial application, secondly not consist of excluded subject matter per s. 1(2) Patents Act 1977, and thirdly be only infringed by the implementation of a product or process, it is difficult to foresee a situation where an AI alone (without the input of an individual or corporation) would be capable of infringing a valid patent.

Nevertheless if it were so capable and this were not foreseeable by a human, this should not automatically absolve the person 'making arrangements' for the AI from being responsible for

controlling the output of said algorithm, including an infringing act (see also our answer to question 5 regarding culpability). Hence for example it should be possible to obtain an injunction against this person forcing them to stop operation of the AI system.

Meanwhile, depending on the facts it may be for example that any damages may be reduced where the specific and unpredicted output is at issue.

In this regard, liability and damages are naturally an area where the facts of the case are important, and in particular the degree of control of the person 'making arrangements' over the AI system; for example a person training a domain specific AI may be argued to have more control and better foresee how it might implement a process than a person framing a problem to be solved by an existing open-domain AI.

The issues may therefore be analogous to those relating to secondary infringement, where for example the intent of the person providing or enabling the process, and/or how easily they could foresee the recipient using that process in an infringing way, have been determinants in such infringement decisions.

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

As with question 15, the AI itself will only implement an internal process, so infringement by the AI itself is only likely to be to that extent; meanwhile infringement of other steps in a wider process (e.g. in relation to training or input data, or subsequent use of results) are a separate matter.

Clearly it is possible that on some occasions it may be difficult to determine if an AI's internal process infringes a patent. However it will be appreciated that this is also true for example of compiled code, or heavily distributed computing methods.

Hence as with all such process patents, it may be difficult to establish infringement where the process being implemented is not public and cannot be reverse engineered. This situation may be exacerbated by a machine learning AI. However, in those circumstances, robust civil procedure and evidential rules which permit disclosure and experiments (as is already the case in the UK) may be sufficient. The additional disclosure or experimentation required may however increase the costs of litigation.

Conclusion:

It is our position that an AI is any technology whose output or functionality is at least in part a consequence of training rather than programming, and that as such AIs are tools for discovering new and non-obvious things already latent within the domain of their training set and inputs. Legal persons may then patent the industrial applications of these discoveries.

This approach renders moot many of the issues relating to inventorship and ownership raised by the questions above.

However, we also note that AIs are a new and unique form of computing (as is tacitly acknowledged by the existence of this consultation), and their interaction with the statutory exclusions should be treated

with care and, in particular, in a different manner to the approach taken with traditional computer programs.

For any queries, please contact Doug Ealey[redacted] on behalf of D Young & Co LLP.