

Annex - Response form

After you have read the consultation document, please consider the questions below. There is no expectation or requirement that all questions are completed. You are welcome to only answer the questions that are relevant to you, your business or organisation.

A copy of this response form is available to download from GOV.uk.

There are two sections on this form:

- A. Questions arising from this consultation
- B. Information about you, your business or organisation

When you are ready to submit your response, please email this form and any other supporting documentation to Alcallforviews@ipo.gov.uk.

The closing date for responses is at 23:45 on 7 January 2022.

The options for computer generated works, text and data mining and patent inventorship are summarised in the following tables.

We have summarised our preferences below using a traffic light scheme. Our reasoning is then provided in the relevant answers.

Green = Agree / Preferred

Yellow = Neutral / Not Preferred

Orange = (Strongly) Disagree

Computer generated works	
Option 0	Make no legal change
Option 1	Remove protection for computer-generated works
Option 2	Replace the current protection with a new right of reduced scope/duration

Text and Data Mining (TDM)	
Option 0	Make no legal change
Option 1	Improve licensing environment for the purposes of TDM
Option 2	Extend the existing TDM exception to cover commercial research and databases
Option 3	Adopt a TDM exception for any use, with a rights holder opt-out
Option 4	Adopt a TDM exception for any use, which does not allow rights holders to opt out

Patent Inventorship	
Option 0	Make no legal change
Option 1	"Inventor" expanded to include humans responsible for an AI system which devises inventions
Option 2	Allow patent applications to identify AI as inventor
Option 3	Protect AI-devised inventions through a new type of protection

Preliminary point

Scope of the term ‘AI’

We assume that the consultation is based on AIs with a similar order of capability to those currently available; i.e. not ‘strong’ AIs that could be considered conscious or generally intelligent.

We propose to interpret the term ‘AI’ as meaning:

“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.

We anticipate that it may become necessary to further review issues regarding AI and IP in the future, if and when strong AIs that could be considered conscious or generally intelligent become widely available.

Section A

Copyright – computer generated works (CGW)

1. Do you currently rely on the computer-generated works provision? If so, please provide details of the types of works, the value of any rights you license and how the provision benefits your business. What approach do you take in territories that do not offer copyright protection for computer-generated works?

N/A

2. Please rank these options in order of preference (most to least preferred) and explain why.

As noted above, we rank them 0, 2, 1 for the following reasons:

Option 0 – make no legal change: Agree / Preferred

Section 9(3) of the Copyright, Designs and Patents Act 1988 (CDPA), accords authorship to the person ‘by whom the arrangements necessary for the creation of the work are undertaken’. This appears to be a reasonable approach and treats the AI as a tool like any other computer based art tool whose programming contributes to the overall aesthetic of the final piece. An AI may be a more sophisticated (or perhaps just more opaque) tool, but is a tool nonetheless. Even once AIs advance to be more autonomous, there would still be “arrangements necessary for the creation of the work”, and this definition would therefore appear to render the current provisions fit for purpose.

Option 2 - Replace the current protection with a new right of reduced scope/duration: Neutral / Not Preferred

We believe that this option may be an acceptable middle ground; one may assume that an AI has no inherent artistic capability, and/or that it is not entitled for example to moral rights or rights that reflect a personal investment in the artistic properties of the work. As such it may be treated as a design engine, whose output is commercialised by another party, and so the rights afforded to its output may be more suitably similar to those of a registered design, e.g. in the order of 25 years.

We note that, as acknowledged in the consultation itself, typically the dissemination of a computer generated work by the owner of the computer will itself attract copyright protection that greatly exceeds the existing 50 year term available to the person running the computer. As a result, any changes to the duration accorded directly to that person are likely to be ineffectual unless they also percolate downstream to other uses of the work (e.g. dissemination by the owner of the computer).

If the duration of protection afforded to a computer generated work was reduced without alteration of the downstream rights, an unnecessarily complex scenario is likely to develop where a piece of computer aided art comprises within itself a mix of rights of different durations (for example where there are special effects in a movie). This could lead to significant enforcement issues.

Hence this option appears to suggest either a relatively ineffectual change to the duration of the right, to reflect the diminished artistic stature of the source, or the creation of complex works comprising parallel durations. This would appear to be the case unless corresponding changes to reduce the duration of protection were allowed to pass through other phases of the creative and commercial processes of making the art accessible. This would appear to represent a more fundamental, and far reaching, change to copyright protection.

Please also note that (as outlined in the answer to question 3 below), it is possible that computer generated works are not always solely computer generated, or that this may even depend on the type of AI being used, making the law difficult to define.

Option 1 - Remove protection for computer-generated works: Would Strongly Disagree

This would seem to be a disincentive to innovation, and also seemingly an unjustifiable punishment inflicted because of an apparent intractability with the law rather than because of any lack of artistic merit in the resulting works.

There is also a significant problem with where the line is drawn for a 'computer-generated work' – it will be appreciated that many artworks and indeed photographs are generated computationally – not just in terms of how digital brushstrokes blend within an image, or in terms of how a model is airbrushed or warped by algorithms, but for example in terms of how photographs are captured with HDR computation, aperture selection, exposure compensation, colour saturation, focal positioning, and many other elements of composition that would previously have been used in support of according a photographer a copyright now being ceded to a computer. For example, many cameras and phones will change settings depending on whether they detect a face, food, foliage, or a sunset.

In addition, as explained in more detail in response to question 3 below, there is a significant question about whether some computer generated works, and in particular the AI based ones prompting the present consultation, are genuinely solely computer generated works.

It does not seem justified to remove protection for a certain category of artistic works merely because new developments in technology make it more difficult to determine the appropriate level of protection.

In short, this approach runs the risk of unintended collateral damage or at best significant new uncertainties, for no clearly justifiable benefit. We would strongly disagree with this approach.

3. If we introduce a related right for computer-generated works, as per option 2, what scope and term of protection do you think it should have? Please explain how you think this scope and term is justified in terms of encouraging investment in AI-generated works and technology.

As noted above, a right similar to a registered design right of about 25 years would seem appropriate if one considers an AI to have no personal investment in a produced work, and any use and dissemination of it to be inherently a commercial exploitation. Even if the commercial exploitation is to sell it as an artwork, it will be appreciated that an AI can produce an endless stream of such works without tiring (or without new inspiration) by randomly exploring the training space its internal weights represent.

It should be noted that there is a question of whether an AI is the sole author of a work; see for example the snail harps and avocado chairs generated by GPT-3: <https://www.technologyreview.com/2021/01/05/1015754/avocado-armchair-future-ai-openai-deep-learning-nlp-gpt3-computer-vision-common-sense/>

The AI would not produce these works if the user had not provided the necessary prompt. There is therefore a question of whether the user is a co-author of the work, or whether they simply provided the system with a design brief; for human artists, the latter may be true, but for an AI where the input directly influences the

AI's search within its state-space, this is less clear-cut; the phrasing of the same basic request to GPT-3 would result in different works being produced.

This is not just a function of the prompt; clearly also the state-space itself affects the output, and this is a function of the selection and sequencing of training and target data. This may be seen as being further removed from the artwork, but nevertheless sets the bounds of capability and expectation of the AI when producing its works, and frequently specifically defines the intended goal, which the AI populates in a relatively scattergun fashion.

Hence whilst a right for a computer generated work may be appropriate for some works, like a weather map that is deterministically derived from satellite data, by contrast for any AI it will have been intentionally influenced toward its output at some stage by a directing mind; as such it is questionable whether such works can be called solely computer generated. This would lead to the issue of needing to decide whether a particular work was a computer-generated work entitled to protection under the new right, or an artistic work produced using a computer which would be entitled to traditional copyright protection. This issue is likely to become increasingly blurred as AI technology advances.

4. What are your views of the implications of the policy options and of AI technology for the designs system?

The registered designs act has a similar provision to s9(3) CDPA, according rights to the person by whom the arrangements necessary for the creation of the work are undertaken. As such it already sits in a position similar to option 2, but without the associated problems of attracting different durations of protection within a single work. As such, Options 0 and 2 seem to have limited impact on designs.

Meanwhile, for the reasons given previously, removing design rights for computer generated works appears to raise a host of economic, legal and philosophical problems for little or no clear benefit.

5. For each option, what are your views on the risk that AI generated works may be falsely attributed to a person?

For Option 0, there is limited risk; 50 years (plus any additional period relating to the commercial use of the work, such as printing or broadcast) leaves little incentive to risk an accusation of bad faith not just within the legitimate period of protection but also in the relatively distant future.

For Option 2, there is an increased risk simply because there is a proposal to reduce the initial term. The risk may increase if the term also impacts downstream terms such as for printing or broadcasting the work. Hence in this case an explicit provision that a *knowingly* false attribution of the author will result in a loss of rights may be appropriate.

For Option 1, false attribution becomes the only means to obtain rights and so is a likely outcome that could effectively increase the rights duration of such a work, as noted in the consultation itself.

Copyright – text and data mining (TDM)

6. If you license works for TDM, or purchase such licences, can you provide information on the costs and benefits of these? For example, availability, price-point, whether additional services are included or available, number and types of works covered by the licence etc.

N/A

7. Is there a specific approach the government should adopt in relation to licensing?

There should be an exception for orphan works, if there isn't already.

There should perhaps also be clear defaults for when to reasonably assume a work is an orphan work; for example if the apparent owner of a work (e.g. of a webpage) is emailed, but no reply is received within a predetermined period, then the work may be considered orphan for the purposes of training. This enables checks (possibly automated checks) at the volume required for big data training.

There could perhaps also be a threshold for how to handle when a work subsequently turns out not to be an orphan; for example if a work contributes less than N% of the training set (e.g. 1% or 0.1%), then its contribution is both marginal and onerous to remove by retraining and so should be exempt. Meanwhile if it is above this threshold and the owner does not license, then the AI could be retrained without the contributing material. In this case there could also be a 'free to continue' clause so that where the AI is already distributed, the distributed version is not changed as this might change the operation of the AI for downstream users, but new distributions must use the retrained version.

To avoid people abusing this by just relying on everyone's contributions being too small to trigger this it, one could also have a clause that a cumulative threshold applies. If a certain proportion or number of owners refuse to licence, or it can be shown that a threshold proportion of owners would not have had the opportunity to consent / license (e.g. because their cumulative works were provided for this new purpose without clear consent back to the source, as may be the case when mining social media posts, medical records and the like), then the same provisions are triggered.

8. Please rank the options in order of preference (most to least preferred) and explain why.

As noted above, we rank them 2, 0, 1, 3, 4 for the following reasons:

Option 2 - Extend the existing TDM exception to cover commercial research and databases: Agree / Preferred

The present framework allows non-commercial research. However the findings of such research may often subsequently be used for, or inform, commercial purposes.

By contrast, particularly for an AI, it is inherent in the operation of an AI that it embodies at least a partial abstraction of the data it has been trained upon. Consequently, where an AI is the result of any non-commercial research phase of developing such an AI (or its training set), it is not then possible to separate the AI from the source material in a sense that safely makes the exploitation of the AI a subsequent and separate commercial act.

It should also be appreciated that many 'big data' AIs ingest huge quantities of data, and it will be frequently impractical to determine the copyright status of every training item.

Therefore, to enable the commercialisation of AIs that are trained on training sets that include copyrighted works, all else being equal within the existing framework, a further provision to allow AIs to be commercialised appears essential.

Notably however this should not remove from the persons responsible for the AI any liability for copyright infringement by the results of the AI; an AI should not be used to 'wash' the copyright from an earlier work.

This will likely require checking new works for effective plagiarism before commercialisation, as would be necessary for any other commissioned work. Given recent decisions in the music world, the idea of AIs that are 'influenced' by copyrighted works is an interesting future issue, but perhaps also one where the degree of influence could be quantified through experimentation and a review of the training sets.

Hence on balance we believe that the commercial exploitation of AIs that have used copyrighted training data is a necessity, but that any potential associated liabilities can be readily accommodated.

Option 0 – Make no legal change: Neutral / Not Preferred

Making no change would appear to limit the usability of AIs, whilst also failing to reflect commercial realities, viz., that text and data mining clearly are used, in the UK and on UK citizens, e.g. for the purposes of advertising. In these cases typically there is a legitimate right to use the text via terms and conditions as per the current framework, but targeted advertising is clearly a commercial use.

Option 1 - Improve licensing environment for the purposes of TDM: Neutral / Not Preferred

As noted above, many 'big data' AIs ingest huge quantities of data; consequently, the problem is not with licensing per se but with volume; improved licensing would not be of help if it still required identifying the licensor of every webpage in two given languages when training a translation AI, for example.

Hence whilst an improved licensing environment of the type indicated in the consultation is no bad thing, it is not a solution for AIs by itself.

Option 3 - Adopt a TDM exception for any use, with a rights holder opt-out : Neutral / Not Preferred

'Any use' invites unintended consequences. However, it may be closer to the scenario we envisage for option 2. The caveats regarding liability for the output of an AI would be essential.

Option 4 - Adopt a TDM exception for any use, which does not allow rights holders to opt out: Would Disagree

Like option 1, there is a problem with scale; if there was an implicit obligation in a rights-holder opt-out to seek out the rights holder before using training materials, then effectively this would make big training sets unworkable.

9. If you have experience of the EU exception with opt out for rights holders, how has this affected you?

N/A

10. How would any of the exception options positively or negatively affect you? Please quantify this if possible.

N/A

Patents

11. Please rank these options in order of preference (most to least preferred) and explain why?

As noted above, we rank them 1, 0, 2 for the following reasons:

Option 1 - "Inventor" expanded to include humans responsible for an AI system which devises inventions: Agree / Preferred

Our position is that for the foreseeable future, AIs cannot be considered to be inventors.

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

Typically, the planning, forethought, and foresight are provided by the parties who curate a training set for domain specific AIs, select a specific AI architecture, and/or construct/convert a problem into a solvable form by an open-domain AI (like GPT-3 mentioned above).

The AI then generates an output or function in a deterministic manner that may nevertheless be new and non-obvious.

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

Hence our position is 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 from 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 actual output is a deterministic process run on a computer simulating the AI (even if this deliberately includes random seeds or the like to help access different regions of the latent space occupied by the AI). There is no room for 'devising an invention' within this framework.

We would strongly assert that the outputs or functions of the AI are thus mechanistic 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. The person providing the industrial application of the discovery is then the inventor, having used as the AI as a tool to make the initial discovery.

However, this approach alone is not ideal as it places the emphasis of protection upon industrial application alone, and not on the new and enabling feature generated by the AI. It also leaves scope for some AI functions or outputs to fall through the gaps, if they are not subsequently given an industrial application by a human. This could disincentivise the use of AIs for innovation. However, this should be regarded as a minimal issue, since any discovery which does not have an industrial application would not be regarded as patentable under the UKPA.

A further concern of this approach is the issue of incomplete or complex inventorship, particularly where different parties create and exploit the AI. In the situation where a first party had created and trained an AI leading to a discovery which was then made industrially applicable by a second party, the second party would be considered the inventor, and therefore the owner, of an arising patent. The patent could therefore be owned by a party making a rather minimal input, in terms of innovation, time and resources. However, this situation is akin to that often seen in life sciences where identifying particular gene or pathway is only considered to generate patentable subject-matter once an industrial application is identified, often by a subsequent party.

Therefore to assist with the recognition of those who contributed to the output of the AI and/or its practical exploitation, a new provision similar to section 9(3) of the Copyright, Designs and Patents Act 1988 (CDPA), where authorship is accorded to the person 'by whom the arrangements necessary for the creation of the work are undertaken', seems to be a clear and straightforward option, and would unify the approach between copyright, designs and patents. This also appears to be the approach alluded to in the consultation document itself.

It would mean that there is a path to ownership that does not rely on immediately identifying an industrial exploitation (arguably a more restrictive requirement than mere industrial applicability).

As a coda, we would redefine this option as: ““Inventor” expanded to include humans responsible for an AI system which ~~devises inventions~~ generates patentable subject matter”. This would appear to be more in line with the current principles for inventorship, and therefore ownership, of patents based on discoveries.

Option 0 - : Make no legal change Neutral / Not Preferred

The current status quo is acceptable but less than ideal. Currently the practical exploitation of the functionality or outputs of AIs is patentable if it conveys a

concrete technical effect, with the patent being awarded to the person(s) devising the exploitation.

As noted above however this does limit patentability (and the path to ownership) to those who industrially exploit an AI, which does not necessarily coincide for example with those who gave rise to its workings. There appears to be scope in some circumstances for an enforced disconnect between who actually contributed to the patentable result and who it is necessary to name for the purposes of protection. Option 1 provides a means to forestall that possibility.

Option 2 - Allow patent applications to identify AI as inventor: Would Strongly Disagree

As noted previously, our position is that AIs do not, and cannot, invent.

Rather, they are a means to discover new and non-obvious properties that are latent within a space occupied by the AI's training set and inputs, and in this sense an AI may discover but not invent a new thing.

Hence at a purely philosophical level, we disagree.

We also disagree at a legal level, noting again the requirement in section 7(3) to be the deviser of the invention, and similarly the caveat in section 43(3) that someone who merely assists with an invention is not an inventor.

If a machine learning system is currently classed as a mathematical method per se, then it is a sophisticated processor used to assist in the above mentioned exploration of the state space defined by the curated training set and the provided inputs. It is not an inventor.

There are also other legal problems that arise from adding a different class of possible inventor. For example, the focus of inventive step often changes during prosecution in response to cited prior art. If one claims the use of an AI to perform a technical task, this may be reasonably considered to be invented by the person who used the AI. However, if the prior art shows a similar use has been disclosed, then the inventive step in the application may rely instead on a difference in behaviours or capability of the respective AIs in the application and prior art – has this now become an AI invented claim? What does this mean for any ownership or assignments, or the validity of a priority claim?

We would also disagree at numerous other practical levels. Once the AI is the inventor, it is not clear how can it assign ownership of the invention to another party in a manner that can be relied upon not just within patent law but also other law, such as for the securing of loans against IP rights, or licensing or assignment. Clearly it will also fall foul of overseas laws such as the assignment and declaration requirements of the US. Finally where the AI is in the cloud, there may also be a question of nationality where some countries require national first filings or have other citizenship based requirements.

We must thus also consider the implications that any changes to UK law may have on the ability of companies to obtain worldwide patent protection. If it is decided that AIs can be listed as inventors on UK patent applications, this has implications for ownership. If an AI cannot be acknowledged as the owner, or inventor, of a patent application in other jurisdictions, this has implications for applications claiming priority to a UK application which lists an AI as an inventor.

AIs are generally considered to be resource hungry, and are therefore often developed, owned and trained by large companies. If we allow AIs to be listed as inventors on patent applications, large companies in both the tech and pharma space could use AIs to generate large numbers of inventions, and subsequently patent applications. This would be disproportionate to the ability of smaller, people-centred companies and has the potential to stifle human innovation in a field.

In short, defining an AI as an inventor is an attempt to solve a problem that does not exist, and in so doing creates a swath of genuine problems elsewhere.

Finally, the recent DABUS decision does not point towards this approach either; the Court of Appeal upheld that the language of Section 7 and other provisions of the UK Patents Act 1977 could not be interpreted as allowing anything other than a person to be considered an inventor.

All three judges agreed that inventors must be natural persons and thus an inventor could not be an AI. In his dissent, Birss LJ still did not agree that the AI was the actual deviser of the inventions, which he again indicated was not consistent with UK law. Rather, his dissent related specifically to section 13 and whether Dr Thaler had fulfilled the requirement to identify the inventor, arguing that this had been fulfilled in good faith.

We therefore conclude that allowing AIs to be identified as inventors has no support in fact or law and would cause a large number of subsequent problems for any corresponding patents / applications.

Option 3 - Protect AI-devised inventions through a new type of protection: Would Strongly Disagree

There is no need for a sui generis right for AI inventions if one accepts that AIs do not invent.

However, in view of the myriad of problems associated with listing AIs as inventors on patent applications, and discussed above, we consider this option to be slightly more preferable than option 2.

As noted above, AIs are generally deterministic and predictable systems that combine the influence of weights honed by training data to produce outputs that are typically a combination of that training and the current input. Whilst it may be impractical to trace the contributions, this does not alter the fact that it is inherent in the underlying algorithms of the AI that this is the case.

Many discussions of AIs discuss these tools as though they are all identical, which of course is not the case. Creating a new right for AI generated inventions would lead to difficulties in establishing the definition of an AI, and when a tool becomes an AI. In the rapidly developing field of AI, this is likely to be an ongoing issue.

Creating a sui generis right for alleged AI inventions would ghetto-ise them, and irrespective of whether the new right was as comprehensive as a patent right, would also mean that such inventions would be reliant on other states recognising them for the purposes of priority for corresponding protection. If a corresponding sui generis right was required to claim priority then for many countries such AI inventions may never be transferrable. This could have large implications to the pharma and life sciences industries where worldwide patent protection is routinely sought. Meanwhile if other countries accepted that such rights could be converted to local patent rights it would beg the question of the sui generis right's purpose here.

There is also an issue if an invention is considered to have been co-invented by a human and an AI. It is not clear whether such an invention should be protected by a patent or by a new sui generis right. Consideration would also need to be given to whether an invention could be transferred between the patent system and the new sui generis right, for example if the claims were limited during prosecution to exclude the contribution made by the human or by the AI.

12. Would the changes proposed under Options 1, 2 and 3 have any consequential effects on the patent system, for example on other patentability criteria?

Option 1 would have only a small impact, serving to expand the accessibility to ownership of novel and non-obvious AI generated functions and outputs in a more equitable manner than the current effective requirement to be the person who provides an industrial application. The main effect may be to better protect all the contributors to the application. As such, a person who expended considerable resources devising and teaching an AI would be entitled to be considered a contributor to the invention, which is likely to incentivise innovation.

Option 2, as stated in relation to question 11, would have wide ranging impacts, either in terms of requiring downstream changes, or in terms of hamstringing patent rights in the face of a lack of changes to accommodate apparent ownership by a non-person.

The most obvious requirement may therefore be to create a new provision identifying the owners of the AI as the de facto owners of the patent, in a similar way to employers owning employee inventions made in the course of their normal duties.

However, as noted in relation to question 11, whilst the UK may make this accommodation there is no guarantee it would be available elsewhere, making it

difficult to claim priority from UK applications in order to obtain protection overseas.

In addition, declaring AIs to be inventors also opens up questions of inventive step and who the skilled person still is. The skilled human is assumed to have full knowledge, but a limited capacity for inventive step. That is well understood. Meanwhile an AI has a degree of knowledge limited by both its training and its architecture, but within that may have the potential to explore the full space these define; as such it may have limited knowledge but full inventiveness within that. Exactly what training can be assumed, what architectures and capabilities of the AI can be assumed, and what inputs or circumstances can be assumed to generate what inventive capabilities, are all unknowns that create a new degree of uncertainty that, as we have discussed above, is unnecessary if one accepts that AIs are inherently incapable of invention in both the legal and common senses of the word.

Option 3 would clearly require the creation of a new sui generis right. In addition to determining the extent to which the right provides a similar scope and duration of protection, or potentially disadvantages inventions created by AIs through lesser protection, it would also be necessary to define the boundary at which an invention was no longer merely assisted by an AI (by 'running the numbers' provided by the people who created the AI) and instead could be truly said to have invented the invention.

As such the new right would require a clear definition of when an AI assisted invention becomes an AI generated invention for the purposes of protection by that right. The issue of inventions co-invented by an AI and a human would also need to be considered as it is not clear to which protection such inventions would be entitled. It would also appear prudent to enable a right to transfer between patents and this sui generis right if its status moved to either side of this divide at any point. If the protections afforded by the different rights are different, this in turn may have consequences for third parties that need to be accommodated, particularly in relation to legal certainty.

Finally, again, whilst this could be accommodated in the UK, there is no guarantee it would be available elsewhere, making it difficult to protect such inventions overseas, particularly if priority to a UK AI sui generis right is required.

For options 1 and 2:

13. If UK patents were to protect AI-devised inventions, how should the inventor be identified, and who should be the patent owner? What effects does this have on incentivising and rewarding AI-devised inventions?

As noted above, the inventor may be difficult to identify using the currently accepted definition of an inventor. If for example the AI is one that regularly consumes new training data, then arguably a current iteration of the AI is not the

same as the AI that previously gave rise to the invention. Hence one criterion may be a requirement for repeatability.

This point may be considered moot on the basis that the rights will have been transferred to a legal entity by this stage, but again may fall foul of other laws that may treat the absence of a functional inventor as problematic and akin to the inventor's death.

Similarly for large systems run in the cloud, the AI may be ephemeral and may also be trans-national at the time that the invention is made, leading to issues regarding foreign filing licences.

The relationship between the patent owner and the AI can be modelled on the relationship between an employer and employee. If the AI is credited with an invention in isolation then the invention will belong to the legal entity who owns the AI. Typically however one might expect the AI to more commonly be considered a co-inventor, on the basis that an AI trained to distinguish cats and dogs will never invent a new protein folding technique – someone constructed the AI (or the facility for the AI to train itself) and is likely to be considered a co-inventor. It would be strange if the ownership of a patent made by employees and the AI they created or operated went to different parties in the absence of any other clear indication of why this would be so.

14. In considering the differences between options 1 and 2, how important is it that the use of AI to devise inventions is transparent in the patent system?

Again we reiterate our position that AIs cannot in fact devise inventions.

However, if there was a system in place that treated 'AI devised' inventions differently to other inventions, then if the resulting UK or global prospects for such inventions was less than for a conventional invention, there would be a clear incentive to identify the inventor as someone other than the AI. This may lead to issues of false inventorship, and would require the Patent Office to decide more regularly upon issues of inventorship.

15. Would the UK adopting option 2 affect your global patent filing strategy, if so, how?

Yes. As discussed above, changing UK law to allow AIs to be listed as inventors, has implications for anyone wishing to claim priority from a UK application listing an AI as an inventor. For any case that seemed to rely on the functionality or output of an AI and for which the AI would therefore be listed as an inventor, whether this was prima facie an AI invention or not, it would become necessary to first file outside the UK, in a jurisdiction which did not allow AIs to be listed as inventors. This would be required in order to ensure an initial right that could be guaranteed to act as a priority for further overseas applications.

This is in part because, as noted previously, it may not be clear until the prior art is cited whether the inventive step may change focus and come to rely on a particular property of the AI – at which point the apparent status of even a

supposedly 'normal' patent application could shift to being an AI generated invention.

There are also ownership implications associated with listing an AI as an inventor, which would need to be ameliorated, potentially by filing a first application overseas. It will also be appreciated that for a number of years the status of AI inventor applications is likely to be challenged, making a first filing in the UK an unnecessary risk.

For option 3:

16. What term and scope of protection should a new right offer?

If an AI has invented an invention, then the scope should be identical to an invention by anyone else. Otherwise a distinction is being made in the value of the inventive process of an AI. If such a distinction exists then it begs the question of why the AI is being accorded inventorship. In short, the approach should be all or nothing.

However, again this is only limited to the UK. It is inevitable that the prospective scope of global protection of a new right will be less than an existing right, making it immediately unappealing.

There is also a question of whether this would affect our relationship with the EPC, if the EPO granted patents for validation in the UK that the UK would not otherwise recognise as such, perhaps requiring conversion to the new right. This is unlikely to be seen as maintaining convergence with the EPC. We would not want to see any move that undermined the UK's participation in the EPC.

17. What should the criteria for grant of a new right be and why? Particularly should it:

- a) Replicate the current requirements for a patent?
- b) Set a different bar for inventive step?
- c) Be an automatic or registered right?

17a. As per question 16, if an AI is an inventor then the bar for invention should be identical. If not then it is not inventing. This also limits the incentive for inventor misattribution (ignoring the overseas rights issues).

17b. As noted in question 12, an AI may have a different inventive process to a human, if it is decided by fiat that what an AI does is inventive. The problem here is that unlike humans, there are myriad 'species' of AIs with different processes and capabilities. If a different bar for inventive step is set in recognition that the different qualities of AIs necessitate or justify this, then it will open the floodgates to a proliferation of standards of inventive step, or an inevitable inequity if a standard is based on one (currently) popular form of AI.

It may also cause a new front for argument as to whether an invention is currently due to the properties of an AI and hence under what criteria it should be examined. For example, it may be decided that AIs are exempt from knowledge that is offline such as paper-only books, as there is a barrier to the assimilation of their knowledge that is not present for a person. This argument could be used to exclude such a book from consideration as prior art. Conversely, AIs could be considered able to use machine-readable data such as binary or encoded information not normally treated as practically accessible to the human skilled person; such information could be excluded by arguing the other way.

Again, these are issues that only arise if one considers AIs to be capable of devising inventions, which we maintain that they are not.

17c. As with 17a., if an AI is an inventor then the bar for invention should be identical. We infer that this question is considering the provision of a lower right such as a utility patent.

Such a right would appear to acknowledge that there is a qualitative difference in human invention and AI invention, and as such appears to represent a fudge that does not have justifiable basis from either perspective.

Again, as per question 16 the approach has to be all or nothing. If 'all' is unpalatable, then options 0 or preferably 1 are the correct ones to consider.

General

18. What role does the IP system play in the decision of firms to invest in AI?

Without speaking from experience, we assume that AIs are used to solve problems that have not been solved by more conventional means, and this is the primary driver for investment.

Having then got an AI, the approach of a given IP system to AI may then play into the decision of where firms pursue protection for it.

The patent system is designed to reward innovation, and to compensate companies for the resources required for innovation. This applies equally to inventions involving AIs. AIs are resource heavy, and companies should be rewarded for innovations involving the use of AIs in the same way as traditional inventions.

19. Does the first mover advantage and winner-take-all effect prevail in industries adopting AI? How would this affect the impact of the policy options proposed on innovation and competition?

Our conclusion from the above answers is that any move to change UK law in a way that makes protection for AI inventions unique, either in terms of the protection offered or in terms of chain of title, are likely to have a significant negative impact on early adopters, who would be better advised to first file (or in the case of chain of title, perhaps even invent) such cases elsewhere. There are also negative implications to the UK taking any action which could be seen as diverging from the EPC, particularly following Brexit.

20. How does AI adoption by firms affect the economy? Does the use of AI in R&D lead to a higher productivity?

N/A

21. Do the proposed policy options have an impact on civil society organisations? If so, what types of impacts?

N/A

Section B: Respondent information

A: Please give your name (name of individual, business or organisation).

████████████████████ on behalf of D Young & Co LLP

B: Are you responding as an individual, business or on behalf of an organisation?

- 1) Business – please provide the name of your business
- 2) Organisation – please provide the name of the organisation
- 3) Individual – please provide your name

Business: D Young & Co LLP

C: If you are responding on behalf of an organisation, please give a summary of who you represent.

N/A

D: If you are an individual, are you?

- 1) General public
- 2) An academic
- 3) A law professional
- 4) A professional in another sector – please specify
- 5) Other – please specify

N/A

E: If you are responding on behalf of an organisation, are you?

- 1) An academic institution
- 2) An industry body
- 3) A licensing body
- 4) A rights holder organisation
- 5) Any other type of organisation - please specify

N/A

F: If you are responding on behalf of a business or organisation, in which sector(s) do you operate? (choose all that apply)

- 1) Agriculture, forestry and fishing
- 2) Mining and quarrying

- 3) Manufacturing – Pharmaceutical products
- 4) Manufacturing – Computer, electronic and optical products
- 5) Manufacturing – Electrical equipment
- 6) Manufacturing – Transport equipment
- 7) Other manufacturing
- 8) Construction
- 9) Wholesale and retail trade; repair of motor vehicles and motorcycles
- 10) Transportation and storage
- 11) Information and communication – Publishing, audio-visual and broadcasting
- 12) Information and communication – Telecommunication
- 13) Information and communication – IT and another Information Services
- 14) Financial and insurance activities
- 15) Real estate activities
- 16) Scientific and technical activities
- 17) Legal activities
- 18) Administrative and support service activities
- 19) Public administration and defence
- 20) Education
- 21) Human health and social work activities
- 22) Arts, entertainment and recreation
- 23) Other activities – please specify

17. Legal activities – Intellectual Property

G: How many people work for your business or organisation across the UK as a whole?
Please estimate if you are unsure.

- 1) Fewer than 10 people
- 2) 10–49
- 3) 50–249
- 4) 250–999
- 5) 1,000 or more

250 people

H: The Intellectual Property Office may wish to contact you to discuss your response.
Would you be happy to be contacted to discuss your response?

Yes

I: If you are happy to be contacted by the Intellectual Property Office, please provide a contact email address.

[REDACTED]

[REDACTED]

J: Would you like an acknowledgement of receipt of your response? Yes/No

Yes