

The background of the entire page is a complex, abstract network of glowing green and red lines and nodes, resembling a molecular structure or a data network. The lines are thin and interconnected, creating a dense, web-like pattern. The green lines are more prominent on the left side, while the red lines are more prominent on the right side. The overall effect is a sense of dynamic, interconnectedness.

wigg!n

**A Response
to the UKIPO's Open Consultation:**

Artificial intelligence and intellectual property

30 November 2020

A response to the UK Intellectual Property Office’s call for views on Artificial Intelligence (AI)

30 November 2020

Contents

PATENTS.....	2
COPYRIGHT AND RELATED RIGHTS	41
DESIGNS	47
TRADE MARKS.....	56
TRADE SECRETS.....	62

PATENTS

Introduction

To fully consider the practical implications of artificial intelligence (“AI”) for patent law in the UK and to assess the UK patent system’s ability to cope with and adapt to the growth of AI technologies it is necessary to have regard to (i) the current sophistication of AI technology and (ii) the patent landscape relating to AI technologies and how it has changed over time.

In order to provide complete responses to the patent questions set out in this consultation, we sought input from Professor David Barber, Director of the UCL Centre for Artificial Intelligence, Director of the UKRI Centre for Doctoral Training in Foundational AI and Fellow of The Alan Turing Institute. We strongly believe it is critical for legislators to have an accurate understanding of the real capabilities of AI, before considering how best the UK’s legislative framework needs adapting and, unfortunately, such an understanding is not easily discernible from the large amount of publicly available written opinion.

With the assistance of Cipher, a patent analysis and IP business intelligence tool, we have also investigated how patenting of AI has evolved across the world to get a better picture of the patent landscape as it relates to AI technologies.

AI research and innovation – a view of the real science

‘Artificial Intelligence’ is a term that is often understood by the wider public as encompassing any technology that performs large scale data processing and analysis. However, within the field of AI research, the term has a more specific meaning and is used to refer to technologies that attempt to replicate (and potentially exceed) human capabilities, including vision, image and text recognition, natural language processing and speech recognition¹.

AI itself has historically been categorised into two categories: symbolic AI and non-symbolic or sub-symbolic AI.

- **Symbolic AI:** heavily rule based and used to solve very logical problems or prove certain results. It involves a manipulation of symbols, which have explicit meaning and can be directly related to something in the real world e.g. a word, the position of a particular piece on a chess board.
- **Non-symbolic or sub-symbolic AI:** the objects being manipulated do not necessarily have any specific meaning or direct correlation to something in the real world. In a neural network, for example, the process occurring in the computer system manipulates an object that is in some way related to something in the real world but not directly so. What is input into the system can be interpreted in a symbolic form but the process occurring in the system does not directly relate to the input symbols or have a simple interpretation in terms thereof.

Machine learning is a data driven approach to solving an AI problem. It involves obtaining large sets of data and labelling the data to help find relationships and mapping. A machine can then be trained using the training data to learn the mapping of each segment of data to the corresponding labels so that it can then replicate the process itself with new data feeds. This approach eventually dominated over more classical scientific reductionist approaches that were common when AI technologies were in their infancy in around the 1940s². In relation to activities such as game play and speech recognition, researchers found that trying to understand and identify the underlying rules and laws and find simple

¹ Note that no single definition of AI is accepted by all technical experts and legal practitioners. For a discussion on the difficulties associated with defining AI and the need for a legal definition see Gary Lea, “Why we need a legal definition of artificial intelligence”, World Economic Forum (7 September 2015), <https://www.weforum.org/agenda/2015/09/why-we-need-a-legal-definition-of-artificial-intelligence>.

² English mathematician Alan Turing introduced AI as a concept in a 1950 paper, and American computer scientist John McCarthy coined the term “artificial intelligence” during the Dartmouth Conference in 1956.

explanations for human processes proved ineffective when compared with a data driven approach, which does not require any understanding of how the underlying processes work.

AI researchers' focus is typically on the end goal i.e. what human capability the AI has been designed to achieve. Accordingly, while machine learning is a type or application of AI technology, it is arguably better regarded as a mechanism to achieve the AI end goal.

AI has broad applicability (some of the applications are mentioned in the paragraphs above). The applicability of a particular AI system depends predominantly on how you train it. Until the last 5 years or so, most of the research into AI related to making performance systems i.e. systems that are effective at producing a specific result e.g. recognising objects and images or understanding speech. However, more recently there has been a tendency for researchers to be more interested in other functions or applications, such as explainability. AI systems were not originally trained with explainability in mind, but the application of AI systems in the public arena has given rise to a need for AIs to be capable of not only producing a result but also of being able to provide an explanation for why that outcome was reached.

In Professor Barber's view (and many other experts in the field), AI technology is still at the stage where AI is being used by humans as a tool (albeit a complicated tool with various applications). We are still many years away from an AI being 'sentient' or acting independently or indeed an AI whose contribution to a patentable invention is so great that there is no identifiable human contributor. In the DABUS case³, for example, a simple algorithm came up with a new useful shape but it was a human that evaluated the output of the AI and considered that the shape might have utility. That is a clear example of AI being used as a tool. Professor Barber does not consider there to be any likelihood or danger in the immediate or mid-term future of an AI system inventing something and evaluating its utility without there being a human involved in the process.

In Professor Barber's opinion, there also does not seem to be any overwhelming reason or impetus to give an AI legal personality. Given the current performance of AI technology, there must still be a human who (at least) looks at the output of the AI and recognises it as an excellent idea that may be applied toward the making of an invention. Accordingly, such a human should be entitled to be the named inventor and owner of any patent for the invention. An alternative way of viewing this is by considering that the person who directs and controls what material the AI produces, should be entitled to own the creative output of their AI, including any patents in the same.

In terms of who should take responsibility for unlawful acts or wrongs committed by an AI, Professor Barber's view is that there is no reason to consider this any differently from how legal principles apply to individuals using other tools. Much would depend on what was reasonable in the circumstances and would likely need to be considered on a case by case basis.

As regards the need for there to be a sufficient disclosure for the invention to be repeatable, the rise in open source AI has gone some way to resolving this potential issue. Some of the leading examples of AI, including GPT-3, a highly sophisticated natural language processing and generation algorithm developed by the AI research lab, OpenAI, have been disclosed online in a sufficient level of detail that with sufficient resources one could essentially reproduce them (even in circumstances where the underlying source code itself has not been publicly released). There are many blog posts and repositories online that explain the approach taken by the AI developers and the specific architecture that is needed to create the invention. In some cases, the training data sets and algorithms are also made available for download. In this way, a blueprint of the AI is provided that allows someone to follow the process and produce the desired result even in circumstances where it is not possible to explain how the AI itself works.

³ BL O/741/19

AI patent landscape

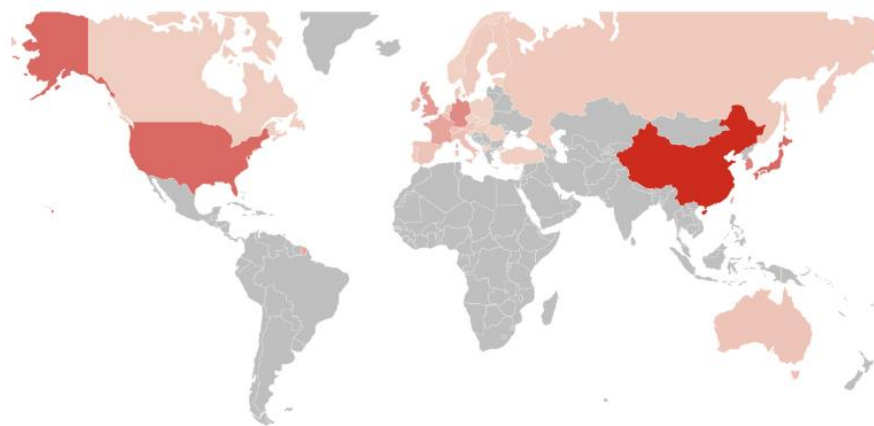
We analysed a report provided to us by CIPHER of all active patent families filed and granted globally since 2010 relating to the field of AI technology. The patents were divided into two categories: (1) those relating to core AI – specifically inventions for the AI itself e.g. computer hardware, nodes in neural networks etc. and (2) those involving the application or use of AI in other fields of technology.⁴

Using CIPHER's online platform (which itself uses AI and supervised machine learning to aggregate and analyse the world's patent data), we were able to identify key trends relating to ownership, geography, fields of application and grant rates of AI patents.

Of the 146,168 currently active patent families (granted or pending) relating to AI technology that were identified in CIPHER's report, 13.2% relate to core AI 'focussed' technology (referred to as 'AI core' in the report) and 86.8% relate to applications of AI technology or AI assisted technology (referred to as 'AI broad' in the report).

Geography

In terms of the territories where patent protection for AI technologies (including both AI core and AI broad) is being sought, it can be seen from the dark red areas on the map and the chart below that the largest proportion of individual granted patents are being obtained in China. The UK appears sixth in the list of countries, with 1.64% of the world's granted AI patents.




CHN: 46%, USA: 16%, JPN: 15%, KOR: 5%, DEU: 2%

Currently active individual patent grants per country. Numbers are aggregated if multiple organisations are selected.



⁴ Note: The results do not include any duplicates, any patent families whose status has been categorised as rejected, expired or inactive or any patents registered in the name of private owners (i.e. individuals not affiliated with a company).




China	24,565
United States	8,385
Japan	8,229
Korea, Republic of	2,944
Germany	1,131
United Kingdom	882
EPO	880
France	749
Taiwan	534
Ireland	363
Next 57	5,130
TOTAL	53,792



Currently active individual patent grants per country. Numbers are aggregated if multiple organisations are selected.

However, given that many Chinese companies only apply for patent protection in China and not elsewhere in the world, it is also useful to consider the data without including China as a territory. The chart below more clearly demonstrates that other key territories where large numbers of AI patents are being granted include the US, Japan and Korea.




United States	8,385
Japan	8,229
Korea, Republic of	2,944
Germany	1,131
United Kingdom	882
EPO	880
France	749
Taiwan	534
Ireland	363
Italy	354
Next 56	4,776
TOTAL	29,227



Currently active individual patent grants per country. Numbers are aggregated if multiple organisations are selected.

Ownership

The top owners globally of AI patent families (granted and pending) are shown in the charts below. 9 of the 10 biggest players are Chinese organisations, with the State Grid Corporation of China coming out at the top of the list. Other important non-Chinese companies include the Japanese multinational, NEC Corp, the US multinational, Intel Corporation, and the Korean multinational, Samsung Electronics.




CIPHER

Currently active patent families (granted or pending) by organisation.

State Grid Corporation of China	3,877
Chinese Academy Of Sciences,Municipal Government of Xiamen City	3,305
Ping An Insurance	2,835
Tencent Holding Ltd	2,640
Baidu Inc	2,225
Univ Zhejiang	1,939
Xidian University	1,788
Tsinghua University	1,766
Univ Electronic Sci & Tech China	1,775
NEC Corp	1,762
Intel Corporation	1,717
Samsung Electronics Co Ltd	1,685
Tianjin University	1,664
NTT	1,614
Alibaba Group Holding Limited	1,602
Fujitsu Limited	1,533
Southeast University	1,475
South China University of Technology	1,401
Google LLC	1,391
Hitachi Ltd	1,342
Univ Zhejiang Technology	1,237
Nanjing University of Posts and Telecommunications	1,182
Univ Beihang	1,166
Canon Inc	1,152
Beijing University of Technology	1,127
Next 9975	107,649
TOTAL	152,869

If patents granted in China are removed (as shown below), the top 25 AI patent owners are more evenly spread around the world, although a number of Chinese companies, including Huawei, Ping An Insurance, Baidu and Tencent Holding still make it into the top 25 list. Only two European companies, Siemens AG and Bosch GmbH, appear in the list, both based in Germany.



CIPHER

Currently active patent families (granted or pending) by organisation.

NEC Corp	1,742
Intel Corporation	1,712
NTT	1,614
Samsung Electronics Co Ltd	1,586
Fujitsu Limited	1,460
Google LLC	1,380
Hitachi Ltd	1,309
Canon Inc	1,135
Toyota Motor Corp	1,005
Mitsubishi Electric Corp	887
Toshiba Corp	877
Sony Corp	790
Korea Electronics and Telecommunications Research Institute (KETRI)	711
Huawei Technologies Co., Ltd.	692
Panasonic Corporation	671
Ping An Insurance	644
Siemens AG	533
Accenture Ltd	454
LG Electronics Inc	450
Bosch (Robert) GmbH	450
Fuji Film Holdings Corp	449
Yahoo! Japan Corp	439
Baidu Inc	416
Tencent Holding Ltd	403
KDDI Corp	396
Next 9975	9,288
TOTAL	31,493

Whilst the majority of the companies that own the largest number of AI patents operate in the tech, electronics and telecoms sectors, companies in other sectors are also developing AI technologies and obtaining patent protection for their inventions. The automotive company Toyota Motor Corp has a large AI patent portfolio relating to self-driving vehicles and the professional services company Accenture Ltd owns over 400 patents relating to various applications of AI technology. Regardless of the sector, it can be seen from the tables below that, with the exception of Intel, all of the biggest patent owners are applying for more patents relating to the application of AI technologies than to core AI.

Including China as a territory:

CIPHER

	State Grid	Chinese Academy	Ping An Insurance	Tencent	Baidu	Univ Zhejiang	Xidian U	Tonghai U	Univ Electronic	NEC	Intel	Samsung Electronics	Tianjin U	NTT	Alibaba	Fujitsu	Southeast U	South China	Google	Hitachi	Univ Zhejiang Tech	Nanjing University of Post-	Univ Easthang	Canon	Beijing University of Tech	Next 6975	TOTAL
AI broad	3,700	2,810	2,432	2,224	1,838	1,688	1,659	1,448	1,540	1,397	552	1,226	1,489	1,431	1,387	1,009	1,328	1,220	1,032	1,198	1,101	1,060	1,037	1,049	976	95,096	132,925
AI core	177	495	403	416	367	251	129	338	235	365	1,165	459	175	183	215	524	149	181	359	144	136	122	129	103	151	12,553	19,944
TOTAL	3,877	3,305	2,835	2,640	2,225	1,939	1,788	1,786	1,775	1,762	1,717	1,685	1,664	1,614	1,602	1,533	1,475	1,401	1,391	1,342	1,237	1,182	1,166	1,152	1,127	107,649	152,869

Currently active patent families (granted or pending) by organisation and technology.

Excluding China as a territory:

CIPHER

	NEC	Intel	NTT	Samsung Electronics	Fujitsu	Google	Hitachi	Canon	Toyota	Mitsubishi Electric	Toshiba	Sony	Korea Electronics and Tea	Huawei Technologies	Panasonic	Ping An Insurance	Siemens	Accenture	LG Electronics	Bosch	Fuji Film	Yahoo! Japan	Baidu	Tencent	KDDI	Next 9975	TOTAL
AI broad	1,379	549	1,431	1,141	945	1,024	1,167	1,033	963	782	749	675	570	494	598	575	432	417	418	343	423	408	343	349	363	7,779	25,350
AI core	363	1,163	183	445	515	356	142	102	42	105	128	115	141	198	73	69	101	37	32	107	26	31	73	54	33	1,509	6,143
TOTAL	1,742	1,712	1,614	1,586	1,460	1,380	1,309	1,135	1,005	887	877	790	711	692	671	644	533	454	450	450	449	439	416	403	396	9,288	31,493

Currently active patent families (granted or pending) by organisation and technology.

When pending applications are removed and granted individual patents are considered the picture is different. The table below shows that the companies that own the most individual granted patents are not the same as those that own the most active patent families (which include granted patents and pending applications). Intel Corporation, Google and Samsung Electronics come out top of this list, and only one Chinese organisation appears in the top 10 list of owners. This table also provides a useful insight into where the biggest players in the field of AI innovation are seeking patent protection. The top owners are focussing on building their patent portfolios in the US and Japan and to a lesser degree Europe, with the exception of the Chinese Academy of Sciences which applies for patent protection almost exclusively in China.

CIPHER

	China	United States	Japan	Korea, Republic of	Germany	United Kingdom	EPO	France	Australia	Taiwan	Next 57	TOTAL
Intel Corporation	299	1,331	146	112	80	108	72	60	1	138	406	2,753
Google LLC	85	1,242	91	74	168	104	81	79	33	45	659	2,661
Samsung Electronics Co Ltd	71	554	83	258	49	56	49	44	1	9	528	1,702
NEC Corp	56	470	769	12	24	21	25	18	2	8	74	1,479
Fujitsu Limited	57	351	716	17	32	34	33	33	1	0	143	1,417
NTT	9	21	1,146	11	16	16	16	16	0	0	95	1,346
Chinese Academy Of Sciences, Municipal Government of Xiamen City	1,244	24	1	0	1	1	1	1	2	0	7	1,282
Toyota Motor Corp	85	222	554	18	86	45	63	50	3	2	140	1,268
Mitsubishi Electric Corp	75	193	616	24	57	33	34	29	0	22	178	1,261
Hitachi Ltd	79	148	713	12	14	14	14	7	3	16	48	1,068
Next 9990	22,828	3,835	3,449	2,437	620	451	492	412	544	295	2,925	38,288
TOTAL	24,888	8,391	8,284	2,975	1,147	883	880	749	590	535	5,203	54,525


Currently active and granted individual patents per country, by organisation.

It is harder to analyse the breakdown of patents across different industry sectors or fields of technology without reference to the industry in which a particular owner operates, as many patents involving the use of AI have broad applicability to a number of different sectors. However, it can be seen from a

review of a selection of patent titles and descriptions, that a large number of AI patents relate to processes such as language recognition, image recognition and data processing.


Granted patents vs pending applications

As regards the split between the number of granted patents versus applications in different territories, it can be seen from the chart below that this varies greatly. In China there are approximately 4 times more pending applications than there are granted patents in the field of AI technology, whilst in the US and Japan the split between granted and pending patents is roughly equal.



Split between applications and granted patents per country. Numbers are aggregated if multiple organisations are selected.

Focussing only on the US and Europe, it can be seen from the table below that in Europe the number of applications greatly outweighs the number of grants, in contrast to the US (which is comparable to Japan)⁵.

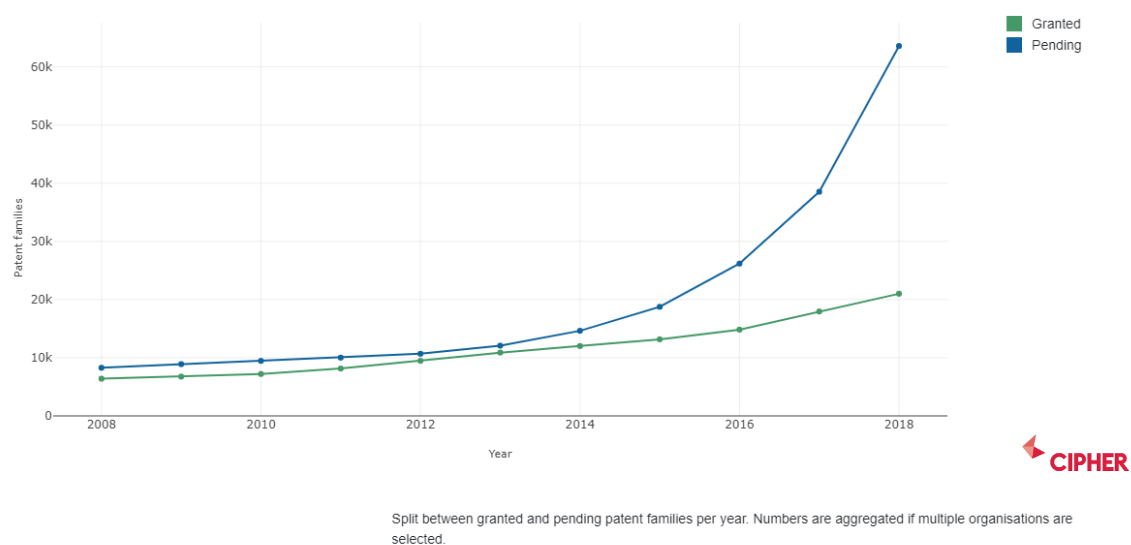
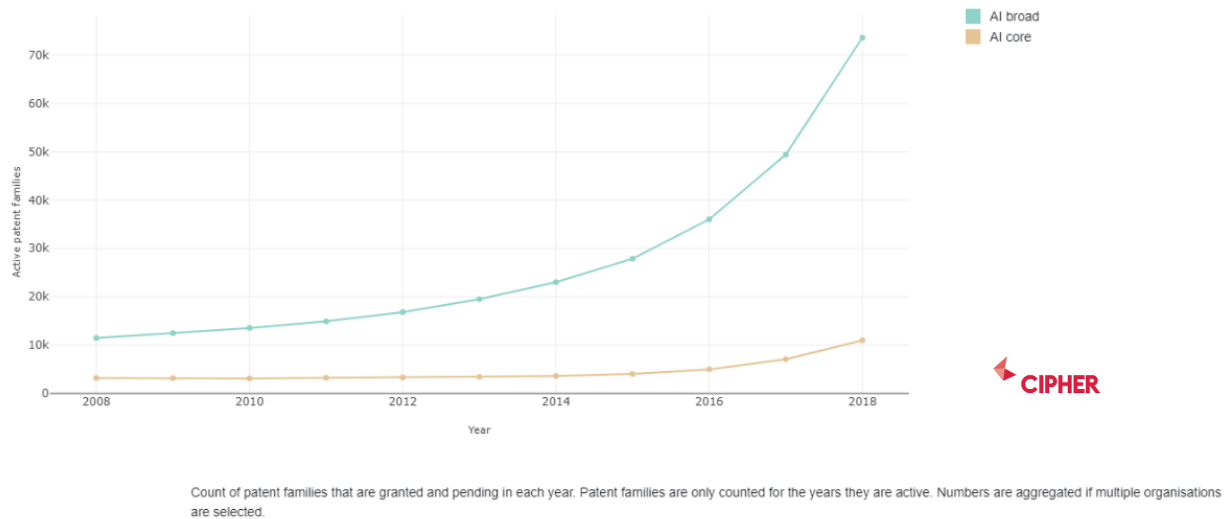


Split between applications and granted patents per country. Numbers are aggregated if multiple organisations are selected.

⁵ It is notable that there are more patents in Germany and the UK than in the EPO. However, this is because the patents in individual countries include both national patents and EP patents designating that country.


Trends over time

Finally, it is worth looking at the trends in the patenting activity over time. It can be seen from the chart below⁶ that over the last 10 years there has been an increase in the number of patent families globally relating to both categories of AI technology and that there has been a dramatic increase in the number of AI patents being applied for in the last 4 years.




From the UK only tables below, the trend is interesting as the first table shows that, in respect of the top 10 owners, there was a large increase in the number of AI patents being applied for between 2014-2018, and the subsequent table shows that the number of granted AI patents has increased in the last couple of years as the pending applications have passed through the system and proceeded to grant.

⁶ Note that the data does not include unpublished applications, so the table excludes figures from 2019 and 2020



Cumulative number of pending and actively prosecuted patent families per year. A family is counted each year it remains pending from priority year.



Cumulative number of granted and active patent families per year. A family is counted each year it remains alive from its grant year.



All of the data set out above demonstrates that AI patents are being applied for both in the UK and around the world in increasing numbers and that patent offices are deeming AI technology to be entitled to receive patent protection, at least to some extent. The question of whether the UK patent system, in particular, is fit for purpose and able to deal in the best possible way with this area of technology in the form that it currently exists, and as it develops in the future, is dealt with in the responses that follow.

Questions

1. **What role can/does the patent system play in encouraging the development and use of AI technologies?**
 - 1.1. The patent system, since its origins, has been designed to further the public policy objective of rewarding and fostering human innovation. This public policy motivation is what underlies the 'patent bargain' whereby '[t]he inventor obtains a monopoly in return for disclosing the invention and dedicating it to the public for use after the monopoly has expired' (see e.g. *Warner-Lambert Company LCC v Generics*⁷). In this way, the patentee secures a return on its investment into research and development of new technologies by having the exclusive right to exploit its invention for a fixed period of time and the public learn of new technological advances, which they are entitled to exploit once the patent monopoly ends.
 - 1.2. This balance between the interests of particular individuals who innovate and invent and of society as a whole is one of the key drivers for inventorship being inextricably linked with first ownership under the

⁷ [2018] UKSC 56, at [17]

patent system in the UK. It is the inventor that is primarily entitled to a grant of a patent and so benefits from its value as a property asset. It is therefore critical for the entity entitled to the patent to be capable of owning and deriving value from its monopoly, otherwise there is no benefit to be had from the 'patent bargain'. For an entity to be entitled to ownership as a benefit it must have a legal personality and without this benefit, the prospect of obtaining patent protection is unlikely to provide any motivation or incentive to innovate. Accordingly, patents that are incapable of being commercialised would appear to serve little purpose from a public policy perspective.

- 1.3. The principle of 'sufficient disclosure' for the purposes of patentability also derives from the patent bargain. A patentable invention requires disclosure of the invention which is 'sufficiently clear and complete for it to be carried out by a person skilled in the art'⁸ or the invention to be disclosed 'clearly enough and completely enough for it to be performed by a person skilled in the art'⁹. Without a sufficient disclosure society cannot take the benefit of the invention because the public is not in a position to understand the steps taken by the inventor and thereby practice the invention on expiry of the patent. In such circumstances, the patent has arguably not been made available to the public in a true sense and the patent bargain has not been fulfilled.
- 1.4. In light of the above, we consider that the patent system in the UK can and does encourage development and use of AI technologies, as it does in relation to any other area of technology, to the extent that the fundamental principles of human inventorship and sufficiency of the invention are able to apply to the AI technology in question so that the patent bargain is maintained. Facilitating the patentability of AI technologies gives rise to other critical issues that require careful consideration and consensus on the appropriate approach, including: what is to be regarded as the state of the art, who the skilled person is, the tests for novelty and obviousness and the test for infringement (each of which is discussed in more detail in the questions below). However, these issues are less likely to have such a major impact on the ability of the patent system to incentivise innovation in AI technologies as the issues of inventorship and sufficiency.
- 1.5. When considering the role the patent system can play in the context of AI technologies, it is helpful to consider the different types of technology and inventions that involve AI. These can be broadly divided into three categories: (1) inventions of which the subject is, or part of it is, the AI itself ("**AI Focussed Inventions**"); (2) inventions involving the application/use of AI, perhaps in another field, e.g. the pharmaceutical sector, which are devised by a person using AI as a tool ("**AI Assisted Inventions**"); and (3) inventions that are devised by an AI ("**AI Generated Inventions**").

AI Focussed Inventions

- 1.6. As regards AI Focussed Inventions, the main obstacle to patentability will likely be whether the subject matter is excluded from patentability, rather than relating to issues of inventorship or sufficiency. This is explained in detail in our response to question 9. There is currently no direct authority on excluded subject matter specific to AI, so guidance must be taken from the law on computer implemented inventions. For, example, for inventions that are directed towards an AI system (including a minimum amount of computer implementation) to be patentable, they must fit around the exclusions contained in section 1(2) PA, which in particular excludes mathematical methods and programs for computers. An AI Focussed Invention (i.e. the AI algorithm) is likely to be characterised as a mathematical method and, as such, may be excluded matter. Accordingly, for the invention to be patentable it must be shown that the AI algorithm produces a non-excluded "technical effect" by its application to a particular field of technology and/or by being adapted to a specific technical implementation.
- 1.7. Issues of sufficiency of disclosure may also arise in relation to AI Focussed Inventions where the AI is used or embodied into the inventive concept. This is because if the disclosure is only 'sufficiently clear

⁸ Art 83 European Patent Convention ("**EPC**")

⁹ section 72(1)(c) UK Patents Act 1977 ("**PA**")

and complete' to an AI, and not a human, the benefit to the public element of the patent bargain is not fulfilled. This is explored in detail in our response to question 11.

AI Assisted Inventions and AI Generated Inventions

- 1.8. In the context of AI Assisted Inventions and AI Generated Inventions, issues of sufficiency may also arise. Where the use of AI would be required to put the invention into practice, sufficient disclosure would be required of the relevant AI functionality and other details such as the information that had been used to train it. As noted in the introductory paragraphs, some of the leading examples of AI are disclosed online in a sufficient level of detail that with sufficient resources one could essentially reproduce them. Conversely, it is unlikely that disclosure of the involvement of an AI entity in devising an invention would be required to satisfy the patentability requirements in circumstances where the use of AI has no bearing on the inventive concept protected by the patent.
- 1.9. Notwithstanding the above comments regarding the potential obstacles to patentability of AI inventions, there is no absolute prohibition on AI-generated inventions being patentable, which we acknowledge would reduce the desire to fund developments in AI technology to the detriment of society. Indeed, as we noted in the introductory paragraphs, it is clear from an analysis of the patent landscape relating to AI technology (including patents to the AI itself and patents to an application of AI in another field) that a large number of patents involving AI are being applied for and granted globally. A considerable number of these are UK and EP patents designating the UK. The data and trends further show that the rate of filing applications for AI patents is increasing and is likely to continue unabated in the short term.
- 1.10. Moreover, we consider that while the patent system in its present-day form does a pretty good job at encouraging and rewarding innovation in AI, other forms of intellectual property right are also worthy of consideration when it comes to providing protection to AI technologies. For example, if an AI invention is not capable of being sufficiently disclosed in a patent specification or if the AI cannot overcome the excluded subject matter restrictions in the PA, trade secret protection may provide sufficient protection instead, especially where the AI itself operates in the opaque environment of its own black box. Where what is sought to be protected is the software in a computer program, copyright protection may be relied upon. In cases where the invention is completely AI-generated it might be appropriate to consider the creation of a new sui generis right, which does not give rise to the same inventor or authorship issues that exist with other existing intellectual property rights. Such a suggestion would need careful consideration at a policy level and be the subject of its own consultation process.

2. Can current AI systems devise inventions?

- 2.1. For AI Assisted Inventions that are made by humans through the use of AI, provided AI is used as a tool for innovation, it should be possible to identify the human(s) who have made the inventive contribution or devised the inventive concept and who are therefore the inventor(s) of the given invention. In this context, the fact that AI is not currently recognised as an inventor does not, in our view, undermine the ability of the UK's legal system to provide effective patent protection.
- 2.2. For AI Generated Inventions with no human contribution to the inventive concept, the UK law does not currently allow for AI to be recognised as the inventor of a patent (the current position on inventorship under UK law is set out in our response to question 3). This gives rise to policy issues as to whether anyone (or anything) should be identified as the inventor and who should be correspondingly entitled to the patent.
- 2.3. However, as we noted in the introductory paragraphs regarding the current state of play in the field of AI, the technology relating to AI is not sufficiently developed such that there is routine and/or widespread use of AIs that are capable of autonomous innovation. The issue of inventorship is therefore more likely to be pertinent to AI in the future as technology develops and AIs truly begin to generate patentable inventions. Further clarity may therefore be needed, as and when this future possibility becomes a reality, in relation to the position of patent applications and, in particular, granted patents that name a

human inventor, but which turn out to have been invented by an AI and lack sufficient human involvement. In any event, we do not consider that the objective of incentivising innovation is furthered by simply allowing an AI to be named as the inventor of a patent because it cannot be the proprietor of the patent (as it is not a legal personality that can hold property) and even if this were not the case, an AI is unlikely to be motivated to innovate by the prospect of obtaining patent protection as it will not derive a benefit from owning or being able to exploit such an asset. These issues are explored in more detail in our responses to questions 3 and 4.

Particularly:

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

- 2.4. In our view, and as corroborated by Professor Barber (see introductory paragraphs), AI technology is still at the stage where AI should properly be regarded as a tool that is used by humans to achieve an end goal. In those circumstances, the end goal has been determined by a human and not the AI machine. We overall agree with the commentators that regard the majority of AI algorithms as not so different from other platform technologies (such as combinatorial chemistry) which are used to produce novel products and processes.
- 2.5. Back in 1988, during a parliamentary debate on the Copyright Designs and Patents Bill, the House of Lords were considering software applications such as computer aided design (CAD) programs in terms of “artificial intelligence”¹⁰. At the time CAD was cutting edge technology and pushed at the boundaries of computer program capabilities, but it was ultimately considered by parliament to be a tool. Looking back on it, this was clearly the right conclusion.
- 2.6. In the context of devising an invention, this conclusion means that an AI system is either the subject of the invention itself (i.e. forming part of an AI Focussed Invention for core AI technology e.g. neural networks) or is being used by a human inventor as a tool to devise an AI Assisted Invention.

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

- 2.7. Under section 130(1) PA, the term ‘inventor’ has the meaning assigned to it by section 7 PA. Section 7(3) PA states that the inventor is “the actual deviser of the invention”. Since no definition of the terms ‘deviser’ and ‘invention’ is provided by the PA, the English courts have considered the terms and provided clarification of their meaning through case law.
- 2.8. In *Yeda Research and Development Co Ltd v Rhône-Poulenc Rorer International Holdings Inc and others*¹¹, the House of Lords held that the ‘actual deviser’ of an invention is the “*natural person who came up with the inventive concept*”. This provides two requirements that must be met in order to establish inventorship. First, the deviser must be a ‘natural person’ and, second, the deviser must have contributed to the inventive concept of the patented invention.
- 2.9. Section 125 PA provides that, unless the context requires otherwise, an ‘invention’ shall be taken to be that specified in a claim of the specification of the patent application or the granted patent as interpreted by the description and any drawings contained in that specification. However, in *Yeda*, the House of Lords held that to be considered a deviser of an invention it is not enough that someone contributed to the claims, because they may include non-patentable integers derived from prior art. Instead, the contribution must be to the formulation of the inventive concept. The determination of inventorship therefore requires a factual assessment of the nature of the contribution to the invention, with only an inventive contribution that goes to the ‘heart’ of the invention being sufficient to confer inventorship. In

¹⁰ <https://hansard.parliament.uk/Commons/1988-04-28/debates/ca8086f9-ed28-446e-ad6b-335e81cdf933/CopyrightDesignsAndPatentsBillLords?highlight=%22artificial%20intelligence%22%20copyright#contribution-84118677-4142-4459-920e-51305a7afe5e>

¹¹ [2007] UKHL 43

contrast, a non-inventive contribution that only amounts to adding to the common general knowledge in the art, such as by contributing only simple and routine experimentation, is not sufficient to confer inventorship¹².

- 2.10. Accordingly, whether a particular individual, such as the AI developer, the user of the AI, or the person who constructs the datasets on which AI is trained is entitled to claim inventorship in relation to an invention made using AI as a tool will depend on the extent of the individual's contribution to the "inventive concept" and whether that contribution can be said to go to the 'heart' of the invention (*Yeda*). For the most part this will turn on the nature of the invention and facts surrounding the specific contribution in each case. We set out below a number of hypothetical scenarios¹³ and provide our view on the likelihood that the individual involved would be entitled to claim inventorship.

1 *A person uses AI to design a particular type of product or process, when the resulting patentable invention made using the AI is of the type of product or process the person intended*

- 2.11. In this scenario, AI is being used by a human to help them invent, but the intent for the result lies with the user. The AI is being used as a tool to 'perfect', assist or refine the human's work product and therefore the contributions by the human are likely to include selecting the field in which to develop the product or process, identifying what they are trying to achieve and how to achieve it, the original work on the invention before the AI is introduced to develop the invention further, and the subsequent use of the AI and its application to the results to support a patentable invention. Such a contribution is likely to be sufficient to establish inventorship.
- 2.12. Whether or not the user intended the results obtained does not directly impact on whether the user could be considered an inventor under UK law as there is no test for intent in relation to inventorship under UK law. However, if the user intended the results obtained it may help to support the contention that his or her contribution goes to the 'heart of the invention'. Although the human contributor(s) may not have envisaged the precise result that the AI produced, the contributor(s) expended intellectual effort, for example, in identifying how to achieve the goal and merely used the AI to optimise the implementation of that goal.
- 2.13. Put another way, the contributor(s) could be considered to have identified the inventive concept from the results produced by the AI, namely the characteristics that led to the patentable invention, and expended intellectual effort to devise a general implementation of a product or process that embodied the inventive concept. AI simply generated the results from a specific implementation of the product or process provided by the human.
- 2.14. In this scenario, the AI can be considered as being used in an equivalent manner as a computer program would be used. The AI tool is only operating within the parameters set by the user of the tool to achieve the result identified by that user. The user is therefore arguably the deviser of the outcome achieved using the tool and so would likely be considered an inventor under UK law.

2 *A person uses AI to achieve a particular intended goal, when the resulting patentable invention made using the AI is not directly related to that intended goal*

- 2.15. As noted in scenario 1 above, there is nothing under UK law that prevents a patentable invention relating to one goal being created in circumstances where the inventor was seeking to address an alternative goal. As such, it is the actual deviser of the means to achieve the goal addressed by the patentable invention who is the inventor, and the patentable invention not being directly related to the originally intended goal does not alter the assessment of inventorship.

¹² *IDA Ltd v University of Southampton* [2006] EWCA Civ 145

¹³ Similar scenarios were considered as part of AIPPI's 2019 Study Question Q272-SGL-2019-en, and we highlight some of the principles explored in the UK Group's response (<https://aippi.soutron.net/Portal/Default/en-GB/RecordView/Index/3806>)

- 2.16. Under UK law, discoveries as such are excluded from patentability¹⁴. In this example, the AI only identified the relationship leading to the patentable invention, but presumably without fully recognising its utility such that it could lead to a patentable invention. It is therefore likely that the human contributor(s) would be recognised as having devised the patentable invention that is in some way based on that useful relationship. This means the human contributor would be considered the actual deviser of the patentable invention, and the AI has again been used merely as a tool.

3 *A person develops or contributes to the design of the AI algorithm that is then used in the making of a patentable invention*

- 2.17. In scenarios 1 and 2 above, the human contributor uses the AI in examples of AI Assisted Inventions. Should the use of the AI algorithm in these examples above fulfil the usual tests for patentability, then there is no reason why a human who has designed or contributed to the design of the AI itself should be excluded from being considered as an inventor of the core AI technology i.e. the AI Focussed Invention, so long as he or she satisfies the test of being an ‘actual deviser’ of the AI Focussed Invention in line with UK case law.

- 2.18. In terms of whether the AI developer can be considered an inventor of a patentable invention that is made using the AI algorithm i.e. an AI Assisted Invention, this person could be considered analogous to a toolmaker. If the “toolmaker” could be said to have (at least jointly) devised the patentable invention, then they could be considered at least a joint inventor of that invention. This would nonetheless still require the “toolmaking” to have contributed to the inventive concept of the invention but this is not outside the scope of an appropriately drafted patent claim. Where the person who designed the AI was unconnected with the specific use that is made of the AI to devise the patentable invention in question, it would be more difficult to establish how they contributed to the inventive concept and thus establish inventorship, in the same way that proprietors of general application software may not claim inventorship in patentable inventions that are created by others using that software.

4 *A person selects or constructs the datasets that are used to train the AI algorithm that is then used in the making of a patentable invention*

- 2.19. In the context of patentability of AI systems, the UKIPO examiners follow the EPO guidelines which state that where an invention is an AI classification method, and that classification method serves a technical purpose (such as classifying digital images or audio based on e.g. pixel attributes or edges), “the steps of generating the training set and training the classifier may also contribute to the technical character of the invention if they support achieving that technical purpose” .

- 2.20. Extrapolating these guidelines to this scenario, it would appear that in the context of an AI used for the purpose described in first two scenarios above, a person who selects training data (and presumably a person who trains the AI) makes a technical contribution to the core AI itself, and as such could be a co-inventor of an AI Focussed Invention in much the same way as a designer of the AI algorithm referred to in scenario 3 above could be.

- 2.21. However, it does not follow that this person would necessarily be a co-inventor of an AI Assisted Invention, as this will be a question of the degree of dependence on the algorithm’s operation on the training data. Returning to the “tool” analogy above, if there is low correlation between the training of the AI and the nature of the results it yields then it is more likely this person would be equivalent to someone who calibrates a tool, and so unlikely to be an “actual deviser” of any invention ultimately produced using the AI.

5 *A person generates or selects the data or source of the data that is input to the trained AI algorithm that is then used in the making of a patentable invention*

¹⁴ section 1(2)(a) PA

- 2.22. This scenario and the one that follows, identify two other types of individual who may be involved in the process of making a patentable invention. Whilst these individuals are not expressly covered by the scope of this question we consider them to be worthy of consideration.
- 2.23. In this scenario, there is no possibility that the person who selects the data to be input into the already trained machine could be regarded as making a technical contribution to the AI Focussed Invention itself, as by this stage the AI algorithm has already been designed and trained.
- 2.24. As to whether this person could be considered to be a co-inventor of an AI Assisted Invention, a person who only generates or selects the “active” data set for the AI (akin to generating or selecting parameters within which software is to operate, which is regarded as “directing” rather than “using” a tool), is unlikely to be an “actual deviser” of any invention ultimately produced by using the AI unless there is a meaningful correlation between the specific selection of the data that is input to the AI machine and the nature of the inventive concept. Each case will turn on its facts. For example, in certain circumstances the person selecting the data may be more likely to be an actual deviser than the person in scenario 4 who selects the training data given that they could be more closely connected to the ultimate invention than someone who merely trains an AI in a hypothetical vacuum, with no knowledge of its final use which then yields an invention.

6 *A person selects one from a large number of outputs produced by the AI used and recognises it to be a patentable invention.*

- 2.25. Whether the selecting human can be an inventor of an AI Assisted Invention, depends on the nature of the inventive concept, and there is a distinction to be made between the AI output itself being a patentable invention and the output simply suggesting to the human that the output could form part of a patentable invention.
- 2.26. In the first case, where the outputs produced by the AI are themselves “inventions”, it is less likely that a person who simply selects one of the outputs and recognises it to be patentable will be an actual deviser of that AI Assisted Invention, in the same way that a person who realises that something created by another person is patentable does not then become the inventor, or co-inventor, of that invention. Similarly, mere selection of something that already exists under UK law is likely to be regarded as similar to a discovery, which falls under the excluded subject matter provisions in the PA. In either case, such a person is unlikely to be an inventor.
- 2.27. However, in the second case, recognising that the discovered output could be applied in a new and inventive way, which could be patentable, may well qualify as a contribution to the inventive concept and allow the human to be classed as an actual deviser under UK law. A person in this situation could be classed as the inventor of a “selection invention” where the outputs produced by the AI exist within the prior art (e.g. where the AI has generated a large number of compounds from a known generic formula and the “invention” is the selection of an individual compound which is not specifically disclosed in the prior art).

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

- 2.28. There are currently very few (if any) situations in which we consider that an invention has been independently devised and created by an AI such that its contribution to a patentable invention is so great that no human contributor to the inventive concept can be identified. As noted in the introduction, in Professor Barber’s expert opinion, we are still many years away from an AI being capable of autonomous invention i.e. an AI machine being able to not only devise a novel product or process but also evaluate its utility without a human being involved in the process. The present-day capabilities of AI systems are such that at least a certain level of human involvement / contribution is required. In these circumstances, a human inventor should in most instances be identifiable as the ‘actual deviser’ of the invention.

- 2.29. Notwithstanding the nature of Dr Stephen Thaler's (the creator and owner of DABUS) arguments before the UKIPO in the DABUS case¹⁵, Professor Barber believes that such AI systems are unable to independently create patentable inventions, in the manner suggested by Dr Thaler. The AI in question used a simple algorithm to produce a new useful shape but it was the human who evaluated the output of the AI and considered that this particular shape might have stackable utility that was the real contributor to the inventive concept. Accordingly, there is an argument that that case was a clear example of AI being used as a tool notwithstanding that the focus was on the issue of inventorship.
- 2.30. In our view, there is an argument that the claims made by Dr Thaler in relation to DABUS should be treated with caution and that it certainly does not provide any proof that AI Generated Inventions are commonplace or yet possible. That does not mean that, as the technology develops in the years and decades to come, such inventions will not become a possibility. However, we do not consider that this future possibility currently warrants the large-scale upheaval of UK patent law that would be required in order to allow an AI to be recognised as an inventor (for this to be possible the inextricable link between inventorship and ownership that is fundamental to the UK patent system would have to be unravelled – see more detail on this point in our response to question 3 below).

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

Current position under UK law

- 3.1. UK patent law does not currently allow an AI entity to be named or recognised as a sole or joint inventor in patent applications in the UK.
- 3.2. In *Yeda*, the House of Lords held that the 'actual deviser' of an invention has to be a "*natural person*", which an AI is not.
- 3.3. The UKIPO's Formalities Manual for patent examiners (the "**Formalities Manual**") provides guidelines regarding the designation of inventors in UK patent applications. Section 3.05 of the Formalities Manual, which was updated in October 2019, states that "*an 'AI Inventor' is not acceptable as this does not identify 'a person' which is required by law*". The guidelines further state that if a patent application is submitted in which an AI system is designated as the inventor in Patent Form 7 (a statement of inventorship used when the applicant is not the inventor or the sole inventor to indicate the derivation of the applicant's right to be granted the patent), the UKIPO's Formalities Examiner will request the filing of a replacement Patent Form 7 in which a person is listed as the inventor instead. If such a replacement statement is not supplied within the prescribed timeframe, the application will be taken to be withdrawn under section 13(2) PA. Consequently, only natural persons can be designated as inventors in the UK and it is therefore not currently possible to designate an AI system as an inventor.
- 3.4. The UKIPO followed these guidelines in the DABUS case mentioned in question 2 above. It refused two UK patent applications¹⁶, in which an AI system called "DABUS" had been designated as the sole inventor of the relevant inventions, on two bases. First, the UKIPO found that since DABUS is not a natural person as required by section 7 PA, it is not possible for the AI system to be regarded as an inventor in the UK patent application (it being the expectation of the law that no one apart from humans can be designated as inventors, which also excludes corporate bodies from being designated as inventors). Secondly, the UKIPO's Hearing Officer noted that even if the refusal to designate DABUS as inventor was wrong, the applicant had not satisfactorily established his entitlement to the invention under section 13(2) PA since "*there appears to be no law that allows for the transfer of ownership of the invention from the inventor to the owner in this case [i.e. Dr Stephen Thaler, the creator and owner of DABUS], as the inventor itself cannot hold property*". This latter point is expanded on in the paragraphs below.

¹⁵ BL O/741/19

¹⁶ GB1816909.4 and GB1818161.0

- 3.5. The European Patent Convention (“**EPC**”) also does not allow for an AI entity to be considered an inventor or co-inventor in a European patent application. The European Patent Office (“**EPO**”) made the same finding in the DABUS case as the UKIPO had but in relation to the two equivalent European patent applications¹⁷, in which the DABUS system had been designated as the sole inventor. The EPO’s decision stated that the EPC does not provide for non-persons (i.e. neither natural nor legal persons) as applicant, inventor or in any other role in the patent grant proceedings and that the EPC only refers to natural persons in the context of inventorship.
- 3.6. It is therefore currently not possible for an AI entity to be an inventor or co-inventor of patent applications in the UK, whether national or applied for as part of a European patent.

Issues with allowing an AI to be recognised as an inventor

- 3.7. As to whether UK patent law *should* allow an AI to be identified as the sole or joint inventor of inventions, our view is that the answer is no, at least not in the context of the UK’s current patent regime.
- 3.8. As was discussed in our response to question 1, under the UK patent system, inventorship and ownership are inextricably linked as it is ownership that underlies the patent bargain. It is the inventor that is primarily entitled to apply for and own a patent and to ultimately benefit from its value as a property asset (section 7(2)(a) PA). Accordingly, an inventor has the right to be named in a granted patent and, if possible, in a published patent application (section 13(1) PA). If the applicant for the patent is not an inventor or is not the sole inventor, the applicant must provide a statement identifying the person(s) whom he believes to be the inventor(s) and indicate the derivation of his or her right to be granted a patent (section 13(2) PA). Equivalent provisions exist in the EPC.
- 3.9. Section 7(2)(b) and (c) PA specify that other than the inventor the only other people who may be entitled to the grant of a patent are: (i) person(s) who, by virtue of a rule / law or an enforceable term of an agreement entered into with the inventor before the making of the invention, were at the time of the making of the invention entitled to the whole of the property in it in the UK; or (ii) to the successor(s) in title of the inventor or the person(s) entitled to the property in the patent. Companies that file applications for patents in inventions created by one of their employees are entitled to the grant of the patent under section 39 PA and/or a term of the employee’s employment contract which assigns all rights in an invention to the employer. Conversely, there is no law that would entitle the owner of an AI to apply for the grant of a patent where the AI was the inventor. There also can be no transfer of ownership of the invention from the AI system to the AI’s owner under law or contract since the AI system does not have legal personality so cannot itself hold any property (this was the problem faced by Dr Thaler in the DABUS case). Accordingly, even if the law on inventorship were changed to include non-persons as inventors so as to allow an AI to be identified as the inventor, and the AI were hypothetically able to file the application, the AI would still not be entitled to be the owner of the patent (as it cannot hold property) and no one else would be entitled to apply for it.
- 3.10. In light of the above, without a complete re-writing of the UK Patents Act to separate inventorship from ownership (which would be a huge undertaking) or to include a new provision that entitles, for example, the owner of an AI (or another person who legislators regard as more appropriately entitled e.g. the creator of the AI algorithm, the user of the algorithm, the creator of the dataset used to train the algorithm etc.) to apply for a patent in an invention created by its AI, we do not see how changing just the law on inventorship to allow an AI to be recognised as a sole or joint inventor takes us any further forwards.
- 3.11. An alternative to revising the provisions of the Patents Act as they relate to ownership and entitlement may be to consider whether the law could be changed to consider an AI as a ‘legal personality’ or an ‘artificial person’ so as to get round the issues of an AI not being able to hold or transfer intellectual

¹⁷ EP 18 275 163 and EP 18 275 174

property¹⁸. After all, some commentators have pointed to a company's incorporation as a legal entity as an appropriate analogue¹⁹. We consider this a step too far. Entities such as companies, which are considered to have legal personality, are made up, at their core, of groups of people. Companies have legal duties and obligations, are able to be held liable for wrongful acts, and there is a clear chain of responsibility amongst the people who make up the organisation. Personal accountability is traceable to officers of the company as enshrined in company related legislation and common law principles such as fiduciary duties. No such accountability or chain of responsibility exists in relation to an AI system.

- 3.12. The only benefit of giving an AI legal personality would be to entitle the AI owner to apply for the patent if they could prove they were entitled to own the invention at the time it was made or that there had been an assignment of ownership in the invention from the AI to the AI owner.
- 3.13. Another possibility that has been proposed by industry commentators could be to ignore inventorship entirely and instead create a legal fiction more akin to the employee/employer relationship that focusses on clarifying the entitlement chain for AI inventions or where an AI would otherwise be regarded as a co-inventor. The entitlement to such AI inventions could automatically vest in, for example, the owner of the AI system, without the need to name the AI as an inventor on the Patent Form 7.
- 3.14. The various propositions set out above all have much the same effect, namely that the owner of the AI ends up entitled to the grant of the patent for the AI invention (save that in the last case there would be no named inventor at all if the AI was the sole inventor whereas in the other examples the law on inventorship would have to be broadened from applying only to 'natural persons' to allow an AI to be designated as inventor).
- 3.15. A different approach which may end up with a different person ultimately entitled to the patent in the invention is to introduce a deeming provision into the Patents Act, in a similar way to how the law of copyright has dealt with authorship of computer generated works²⁰. The approach taken in the CDPA could offer a starting point for an equivalent deeming provision for inventions made by an AI without a human inventor, i.e. the inventor could be deemed to be the (natural, and perhaps legal) person who made certain arrangements necessary for the creation of the invention, arrangements which would otherwise be inadequate to establish inventorship. If the person who made the 'arrangements' was not the AI owner, then the AI owner would not be the named inventor or be able to claim entitlement to the patent. However, given the different scope in protection afforded by copyright law, careful consideration should be given to whether such ownership deeming provisions are appropriate for patentable inventions from a policy perspective. Take, for example, an AI based on open source models such that, in so far as the system itself was concerned, the controller of such an AI could not assert itself as the 'owner' at least from a copyright perspective. In the event that such an open source AI system gave rise to an AI Generated Invention, should such a nominal 'owner' also own the inventive output?
- 3.16. Regardless of the mechanism adopted, a number of issues arise. Making it possible to circumvent the current linked principles of inventorship and first ownership creates the potential for monopolising entire areas of technology without having to pay for the effort of a human 'spark'. Even if patents on AI Generated Inventions could be argued to ultimately promote innovation, those patents may "negatively impact future human innovation as supplanting human invention with autonomous algorithms could result in the atrophy of human intelligence"²¹. The concern is that reduced inventive talent could lead to

¹⁸ See Nicolas Petit, "Law and Regulation of Artificial Intelligence and Robots: Conceptual Framework and Normative Implications", at 19 (9 March 2017), at 10-11 and 19: <http://ssrn.com/abstract=2931339>

¹⁹ See Ben Allgrove, "Legal Personality for Artificial Intellects: Pragmatic Solution or Science Fiction?" (June 2004) (Master of Philosophy thesis, University of Oxford): <https://ssrn.com/abstract=926015> for a discussion of legal personality for AI, addressing philosophical personality and legal personality, as well as the different methodologies for defining legal personality.

²⁰ Section 9(3) Copyright, Designs and Patents Act 1988 ("CDPA") provides a deeming provision for authorship of computer-generated works, being works which are "generated by a computer in circumstances such that there is no human author of the work" (section 178 CDPA). The author of a computer-generated copyright work is taken to be "the person by whom the necessary arrangements necessary for the creation of the work are undertaken".

²¹ See Erica Fraser, "Computers as Inventors – Legal and Policy Implications of Artificial Intelligence on Patent Law", SCRIPTed 13(3), 305 (2016) at 327: <https://script-ed.org/wp-content/uploads/2016/12/13-3-fraser-1.pdf?d=11302020>

the elimination of high-quality research and development (R&D) jobs or entire R&D-intensive industries²². Though there are always gaps in various companies' access to funding R&D, a change to make AI Generated Inventions capable of patent protection may further shift the commercialisation opportunities even further in favour of large, well-funded organisations that have exclusive access to expensive AI systems and stifle new ventures by creating barriers to subsequent research²³.

Conclusion

- 3.17. As we have noted previously, whilst we are still at the stage (as we expect to be for some time to come) where in all, or the majority, of cases it will be possible to identify a human inventor of a patent, the current requirements under UK law would seem to be adequate. In those circumstances we do not consider it necessary to implement any of the more far-reaching options that are outlined in the preceding few paragraphs, all of which are fraught with difficulties.
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?**
 - 4.1. As regards AI Assisted Inventions, which, as we have noted throughout our responses, we consider to be the default situation currently, the creation of these inventions will require human input. For example, automated decision making in high frequency algorithmic trading still requires extensive input from humans to define specific parameters and select the correct sets of data. The AI itself is focussed on information extraction and speed of processing. Accordingly, in almost all cases it should be possible to identify the human(s) who have made the inventive contribution and who are therefore the inventor(s) of the given invention. In this context therefore, the fact that AI is not currently recognised as an inventor does not undermine the ability of the UK's legal system to provide effective patent protection or encourage innovation.
 - 4.2. We do acknowledge, however, that even before technology has advanced such that there is routine and/or widespread use of AIs that are capable of autonomous innovation i.e. AIs capable of creating AI Generated Inventions (which, as we have stated, we believe we are still some way off achieving) issues to do with inventorship and AI may arise. It is conceivable that, possibly in the near to mid-term future, AI will become sufficiently advanced to the extent that it will be capable of recommending solutions that humans merely have to test in order to prove that they work. Applying the case law to this scenario, the human(s) will not have met the criteria for inventorship. However, extending the UKIPO's reasoning in the DABUS case, the AI also will not have met the criteria for inventorship.
 - 4.3. In reality, quite opposed to discouraging applicants from applying for patents for such inventions, applicants may instead, when applying for a patent for an AI invention, indicate that they are the inventor or that they derive a right from another human inventor. In such circumstances it will be difficult to show that the human inventor in question has not met the criteria for inventorship. Section 7(4) PA indicates that an applicant who states that they are the inventor or that they derive a right from the inventor is presumed *prima facie* to be entitled to be granted a patent. Para 7.13 of the UKIPO's Manual of Patent Practice (the "**Patent Practice Manual**") indicates that the UKIPO makes no attempt to question these assertions. Correcting the inventorship is therefore only achieved through challenges brought by third parties.
 - 4.4. The Patents Act allows for the identity of the inventor to be challenged through provisions dealing with the mention of the inventor (section 13 PA). As entitlement to apply for and be granted a patent derives from the inventor(s) disputes regarding the identity of the inventor will also be relevant in challenges

²² <https://www.weforum.org/agenda/2017/07/how-long-before-a-robot-takes-your-job-here-s-when-ai-experts-think-it-will-happen> and <http://www.automatinginvention.com/about.html>

²³ See Ryan Abbott, "I Think, Therefore I Invent: Creative Computers and the Future of Patent Law" (28 September 2016) at 1105-1107: http://epubs.surrey.ac.uk/820928/1/01_abbott.pdf

relating to entitlement and ownership (sections 8, 12 and 37 PA). Indeed, commonly UKIPO Patent Decisions dealing with inventorship disputes in the context of sole inventors are brought under both sections 13(3) (mention of the inventor) and section 37/section 8 PA (determining entitlement) to remove the existing name of the incorrectly accredited inventor and to establish the correct inventor's entitlement claim.

- 4.5. A particular feature of entitlement proceedings is that they may only be commenced by another person who claims they are so entitled (sections 8(1), 12(1) and 37(1)). While grant of a patent to a person who was not entitled to the grant is provided as a ground for revocation under section 72(1)(b) PA, section 72(2) similarly limits the right to rely on this ground to a person entitled to bring an entitlement claim under section 37 PA. The net result of this is that it will not be possible to challenge the entitlement of an applicant for/proprietor of a patent in the scenario being considered i.e. an invention devised solely by an AI system but where a human inventor has been named, as the AI system would not itself have standing to challenge the entitlement of an applicant to obtain a patent for its invention as it does not possess its own legal personality and, importantly, third parties would also be precluded from challenging the applicant/proprietor's entitlement under Sections 8, 12 or 37 PA (or the validity of the patent on the basis of entitlement under Section 72 PA) due to lack of standing to do so.
- 4.6. Section 13(3) PA and Rule 10(2) Patent Rules ("PR") appear to provide more scope for third parties to challenge such a patent/application as both allow any person to challenge the mention of inventors in a published patent application or granted patent. Rule 10(2) PR permits a person who alleges that any person ought to have been mentioned as the inventor or joint inventor of an invention to apply to the comptroller for that person to be so mentioned. Section 13(3) also permits any person to make an application for a decision by the comptroller that a named inventor should not have been mentioned as such. Hence, in each case applications under Rule 10(2) PR and s13(3) PA are not limited to persons who claim that they themselves should be named as an inventor. However, applications to remove an inventor under section 13(3) by an uninterested party may encounter difficulties. Any such third party challenge to a person's inventorship would need to overcome the presumption under section 7(4) PA that the applicant is entitled to apply for and be granted the patent. The third party would, for example, need to prove that the named human inventor contributed nothing of substance to the inventive concept or concepts²⁴. A third party is therefore likely to face evidentiary challenges as they are unlikely to have access to detailed information regarding the circumstances under which the invention was made and the interaction between the named human inventor and the AI.
- 4.7. If despite the difficulties explained above the application to challenge a person's inventorship succeeds, this raises the question of what will happen to the patent application or granted patent after that, especially in cases where there are no other named inventors. In respect of patent applications, it is likely that the UKIPO would treat the application as deemed to be withdrawn under section 13(2) PA, following the logic in DABUS. However, there is currently no provision of the PA or PR which would appear to apply in the event the challenge was made in relation to a granted patent. The failure to name an inventor is not, for example, listed as one of the exhaustive grounds for revocation of a granted patent under Section 72(1) PA, but it seems difficult to envisage how such a patent should remain in force where it lacks a named inventor.
- 4.8. If, as we contend, it is concluded that naming an AI as inventor is not desirable, then it could be a relatively straight forward exercise to amend the PA to clarify that patents successfully challenged under section 13(3) to remove the sole inventor are susceptible to revocation where no alternative inventor has been determined. Conversely, if there is agreement and public policy reflected that protection for AI Generated Inventions is desirable, then some of the propositions set out in our response to question 3 might need to be considered. However, as noted in our response to that question, those propositions give rise to a number of other issues which would need to be addressed before any amendments to legislation were made.

²⁴ *University of Southampton's Applications* [2004] EWHC 2107

- 4.9. Notwithstanding the comments above regarding the potential issues that may arise if an AI is not allowed to be named as an inventor and humans are incorrectly named in their place, we do not consider that recognition of an AI entity as an inventor or co-inventor furthers the public policy objective of fostering innovation and may, in certain circumstances, hinder it.
- 4.10. As discussed in our response to question 3 above, amending legislation to only allow for an AI to be considered an inventor (without making other significant changes to rules around entitlement to facilitate a different entity or individual to own the resulting patent) will not encourage an AI owner or developer to seek patent protection for inventions its AI creates since the AI would still not be entitled to be the owner of the patent (as it cannot hold property) and no one else would be entitled to apply for it (as there would be no successor in title to the inventor as the AI can also not transfer property).
- 4.11. Even if the AI were also to be entitled to the grant of the patent (which without giving it legal personality cannot occur), it cannot derive value from its monopoly, it will not be motivated to innovate by the prospect of obtaining patent protection and there is no benefit to be had from the 'patent bargain'.
- 4.12. It could be argued that permitting AI systems to be recognised as inventors and AI Generated Inventions to be patentable would promote disclosure and commercialisation of such inventions on the basis that without such protection, owners of creative computers and AI machines might choose to protect patentable inventions as trade secrets without any public disclosure. However, alternative tools, such as first-mover advantage and social recognition of AIs, as well as alternative technologies that prevent infringement of patent rights, can also lead to innovations and public disclosure of AI inventions.
- 4.13. If it cannot be said that crediting an AI as an inventor furthers the aim of encouraging innovation in AI technology, it equally cannot be said that restricting an AI from being recognised as an inventor stifles innovation.

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

- 5.1. The key reasons for recognising and naming an inventor in a patent application are to: (i) provide a starting point for the determination of ownership of the resulting patent (the inventor will be the first owner of the patent application unless the invention accrues to his or her employer by operation of law, or by operation of an assignment); and (ii) provide the inventor with deserved acknowledgement which can have an impact on their reputation or their employment terms e.g. as part of a company's reward or compensation scheme (in certain circumstances if the invention is of 'outstanding benefit' to the employer, the employee will be entitled to additional compensation²⁵).
- 5.2. Recognition of inventors as part of those persons' moral rights is not considered relevant in relation to AI entities. The AI entity not being a person is not considered to possess any moral rights and there is arguably no moral case for recognising an AI as an inventor. Conversely, in circumstances where an AI is recognised as an inventor to the exclusion of a human who could be said to have contributed to the inventive concept, that person has lost out on their right to first ownership of the patent and their moral right to be recognised for the human 'spark' of invention they have contributed that the patent system is intended to reward.
- 5.3. There may be a case for recognising an AI as an inventor where a natural person employed with one organisation has collaborated with another organisation's AI and had the AI been a natural person it would have qualified as a co-inventor. From a moral perspective it could seem wrong/unfair that the organisation that contributed the AI receives no recognition or ownership stake in any resulting invention. However, this point could be rectified by addressing the potential imbalance in the commercial contract with the AI owning/controlling entity. This point is discussed in our response to question 6 below regarding ownership.

²⁵ *Unilever v Shanks* [2019] UKSC 45

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

- 6.1. As set out above in our response to question 3, the law regarding ownership of patents is inextricably linked with that of inventorship. Section 7(2) PA provides that a patent for an invention may be granted to the inventor or their successor in title. The questions that arise in relation to who should be entitled to the grant of a patent for an AI invention therefore flow directly from who is designated as the inventor.
- 6.2. Under the existing principles of UK patent law, where a person uses an AI in the creation of an invention in circumstances where the invention turns out, at least in part, to be an AI Generated Invention, the natural person co-inventor (or their employer) would receive the sole and exclusive benefit of the patent even though the natural person is only one of the contributors to the inventive concept. This may not pose a practical or moral issue where the natural person (or their employer) has also invented the AI but its implication in a collaboration or joint venture situation is more pronounced. Unlike a situation where the natural person is working with another organisation's employee to devise the invention (in which case that organisation, by ordinary operation of law, would likely have an ownership stake in the patented invention), where a natural person is working with another organisation's AI, the collaborating organisation is deprived of any rights to the invention and the natural person would be the sole inventor and primary owner of the resulting patent.
- 6.3. If it were permissible for the AI to also be named as a co-inventor in the above scenario, then for the naming of the AI as the inventor to have any benefit from an entitlement point of view, the organisation or individual owning the AI (or another entity who is deemed to be equally entitled) should have an ownership stake in the patented invention. As noted in our response to questions 3 and 4 above, this would require further amendment to the current UK law beyond changes to the law on inventorship. A simpler solution could be to provide for more equitable circumstances in the commercial arrangement that supports the collaboration. Indeed, there is likely to be significant bargaining power in the hands of the company who controls the AI that can produce AI Generated Inventions so that it may ask for financial proceeds of commercialisation, thereby compensating it for its investment. Alternatively, the naming of the AI as co-inventor could be regarded as being appropriate purely for information purposes only without there being any consequences for ownership and economic entitlement to the invention in question. Naming the AI as an inventor may require providing information to allow the AI entity to be uniquely identified and might require provision of, for example, the training data provided to the AI entity, which could prove useful in assessing and evaluating patentability – inventive step and/or sufficiency of disclosure. Moreover, recognising the input an AI has had on the inventive concept by naming it on a public register, may facilitate the public disclosure of ways in which AI is being used in society.
- 6.4. Under the current law solely AI Generated Inventions can have no inventor and therefore a patent cannot be granted. If it were permissible to name an AI as an inventor, then in this instance the AI would be the sole inventor. If legislation were adopted to allow AI Generated Inventions to be patentable notwithstanding a lack of inventorship or the AI being named as the sole inventor, then who should benefit? The most obvious persons who could be entitled to own the patent would be the AI owner as the person or entity who has invested in the AI machine and has presumably chosen to use it to innovate. However, as noted in our response to question 3, in a situation where an AI could be designated sole inventor, other individuals may also feel they have the right to claim entitlement to any patent granted, including the creator of the AI algorithm, the user of the algorithm and the creator of the dataset used to train the algorithm. Additionally, as highlighted above, what if large portions of the AI system in question were based on open source licences?
- 6.5. Notwithstanding the hypothetical suggestions above, our overriding view, supported by industry commentators, is that, like any other tool used in the invention process, the AI system could have been an important or even crucial tool in making the invention but is still ultimately used as a tool. The invention is therefore the product of the natural person(s) who devised the inventive concept, not those who provided the tools for doing so. Accordingly, the idea that the natural person(s) who own, designed/programmed the AI system or first saw and appreciated the AI's result deserve recognition (if

they were not the same person(s) who were involved in the creation of the specific invention) has been met with little enthusiasm from industry commentators²⁶.

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

- 7.1. In considering this question it is important to distinguish between 'AI inventions' where the invention is an AI Focussed Invention or incorporates at least the use of a minimum amount of AI in the inventive concept and where AI is used to create an invention, specifically an AI Assisted Invention (but the use of the AI itself does not form part of the inventive concept).
- 7.2. Generally, the use of AI in the inventive process does not by itself lead to the exclusion of patentability of the invention as a whole in the UK. This is evidenced by the large number of patents involving AI technology that are being applied for and granted in the UK (see Introduction).
- 7.3. For an AI Assisted Invention, the involvement of the AI is not a bar to patentability of the resulting invention (e.g. the use of AI in drug discovery is not a bar to the patentability of a pharmaceutical product which results from that drug discovery process).
- 7.4. For an AI Focussed Invention, the patentability of that invention will depend on whether the AI contribution is excluded from patentability under section 1(2) PA. Excluded subject matter and its relevance to AI technology is dealt with in our response to question 9.

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

- 8.1. Beyond resolving the issue of inventorship and overcoming potential issues of excluded subject matter (touched on briefly in our responses to questions 1 and 7 above and dealt with in detail in question 9 below), granting patent protection for AI generated inventions gives rise to other questions. For example, it is unclear whether such patents would require different standards relating to sufficiency of disclosure, inventive step, and who the 'person skilled in the art' should be, and how as a practical matter such applications should be examined.
- 8.2. We consider that the biggest challenges to patentability that may arise as AI technology continues to become more sophisticated are those relating to disclosure and sufficiency i.e. the requirement for a patent to provide enough detail to allow a skilled person to perform the invention. We discuss this requirement in detail in our response to question 11 below. However, we note that this issue will only have an impact in respect of certain types of AI inventions and even in those circumstances, there may be ways to overcome such challenges.

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?

- 9.1. AI inventions may be characterised as an algorithm and so a mathematical method (s.1(2)(a) PA), or a method for performing a mental act or a program for a computer (s.1(2)(c) PA) and as a result may be excluded from patentability to the extent that any such patent relates to those things "as such".
- 9.2. An example of AI at its most complex (currently) is an artificial neural network. Such a neural network may be described as a computing system inspired by biological neural networks comprising a series of algorithms that endeavours to recognise underlying relationships in a set of data through a process that mimics the way the human brain operates. A current example of a sophisticated AI system is a natural language processing application developed by OpenAI, called GPT-3. GPT-3 has undertaken training analysis on a large amount of text data and can build language output, such as sentences, by studying

²⁶ See Ryan Abbott, "Hal the Inventor: Big Data and Its Use by Artificial Intelligence", in Big Data Is Not a Monolith, MIT Press (Cassidy R. Sugimoto et al., eds, 2016) at 12-15: <https://ssrn.com/abstract=2565950>

not just the words and their meanings, but also by gathering an understanding of how the usage of words differs in different contexts. Its algorithms have been trained through unsupervised learning so rather than being amended in response to indications that its conclusions are right or wrong, it instead calculates how likely its output will be what the user needs from the training data itself.²⁷

- 9.3. Such systems are entirely computer-based and irrespective of complexity still require a computer processor to execute lines of software code. As AI inventions are, by definition, computer-implemented, the adequacy of the law concerning the patentability of computer-implemented inventions in the UK is highly relevant. Similarly, under UK law it is not possible to patent a discovery or mathematical algorithm as such, but “if you can tell people how it can be usefully employed, then a patentable invention may result.”²⁸ Therefore, the specific use of algorithms in AI may be patentable, notwithstanding that the algorithms themselves may not.

Case Law

- 9.4. There is a long history of English case law that attempts to clarify the scope of section 1(2) PA. Many of these cases have been used by the UKIPO in formulating guidance for examination. Of most relevance is the decision in *Aerotel v. Telco; Macrossan's Application* [2007] RPC 7, in which the UK Court of Appeal laid down the following 4-step test for the assessment of whether the claimed computer implemented invention is patentable:

1. properly construe the claim;
2. identify the actual contribution;
3. ask whether the identified contribution falls solely within the excluded subject matter; and
4. check whether the actual or alleged contribution is actually technical in nature.

- 9.5. Applying the *Aerotel* guidance in *Symbian Limited v Comptroller General of Patents* [2008] EWCA Civ 1066, the Court of Appeal came to the conclusion that a computer program would not be excluded if it had the knock-on effect of improving the functionality of the computer as a matter of practical reality. Specifically, it accepted that two forms of technical effect may confer patentability on a claimed computer program:

- Where the invention solves a problem within the computer; and
- Where the beneficial consequences of operating the program feed into other devices.

- 9.6. In the *AT&T Knowledge Ventures LP v CVON Innovations* decision²⁹, the UK High Court provided further “signposts” as to what could constitute a “relevant technical effect” in the context of computer implemented inventions, specifically:

1. the claimed technical effect has a technical effect on a process which is carried on outside the computer;
2. the claimed technical effect operates at the level of the architecture of the computer, that is, whether the effect is produced irrespective of the data being processed or the applications being run;
3. the claimed technical effect results in the computer being made to run in a new way;
4. the program makes the computer a better computer in the sense that it runs more efficiently and effectively (broadened from an increase in speed and reliability by *HTC Europe Co Ltd v Apple Inc.*³⁰); and
5. whether the perceived problem is overcome by the claimed invention as opposed to merely being circumvented.

²⁷ “Language Models are Few-Shot Learners”, 22 July 2020, [arXiv:2005.14165v4](https://arxiv.org/abs/2005.14165v4) [cs.CL]

²⁸ *Genentech Inc.'s Patent* [1989] RPC 147 at 208.

²⁹ [2009] EWHC 343

³⁰ [2013] RPC 30

- 9.7. How this body of law should apply to AI inventions to restrict their patentability depends largely on what is meant by an 'AI invention'. To the extent that it is an AI Focussed Invention then it may be said to solve a problem within a computer (*Symbian*) or that by virtue of, for example, the way its 'nodes' are arranged, that it operates at the level of the architecture of the computer or makes the computer run in a new way (*AT&T*). To the extent that it is an AI Assisted Invention, then it may be more akin to the second limb of *Symbian*, where the output feeds other devices and so has an effect on a process that is outside of the computer.

Patent Office Practice

- 9.8. The UKIPO has provided its own guidelines regarding the patentability of inventions relating to AI systems as set out in the UKIPO's Manual of Patent Practice (the "**Patent Practice Manual**"). Paragraph 1.39.3 of the Patent Practice Manual notes that because AI inventions are generally computer-implemented, patentability is typically assessed as any other computer-implemented invention. The Patent Practice Manual highlights types of AI inventions that might or might not fall under the exclusions from patentability under UK law, for example, AI inventions that are directed to a specific technical process outside of a computer or that form part of the internal workings of a computer are not likely to be excluded from patentability while AI inventions that are directed to an excluded process (such as a business method) are likely to be excluded.
- 9.9. In theory, there are parallels to be drawn with the EPO's approach to the patentability of mathematical methods (under which AI is categorised). Part G-II 3.3 of the EPO Guidelines for Examination refers to Technical Applications and Technical Implementations.
- 9.10. Technical Applications that may be patentable include: (a) controlling a specific technical system or process, (b) determining from measurements a required process step in a machine implemented process, (c) digital audio, image or video enhancement or analysis, (d) separation of sources in speech signals and speech recognition, (e) encoding data for reliable and/or efficient transmission or storage, (f) encrypting/decrypting or signing electronic communications, (g) optimising load distribution in a computer network, (h) determining the energy expenditure or the body temperature of a subject or processing data obtained from physiological sensors, (i) providing a genotype estimate based on an analysis of DNA samples, as well as providing a confidence interval for this estimate so as to quantify its reliability, (j) providing a medical diagnosis by an automated system processing physiological measurements, and (k) simulating the behaviour of an adequately defined class of technical items, or specific technical processes under technically relevant conditions.
- 9.11. To the extent therefore that AI-related inventions are challenged as being mathematic methods and so excluded from patentability, should those inventions be directed at applications similar to those above, then they are more likely of being able to satisfy the requirement of technical contribution.
- 9.12. By contrast with Technical Applications, patentable Technical Implementations are described as:

*"a mathematical method that contributes to the technical character of the invention independently of any technical application when the claim is directed to a specific technical implementation of the mathematical method and the mathematical method is particularly adapted for that implementation in that its design is motivated by technical considerations of the **internal** functioning of the computer (T 1358/09). For instance, the adaptation of a polynomial reduction algorithm to exploit word-size shifts matched to the word size of the computer hardware is based on such technical considerations and can contribute to producing the technical effect of an efficient hardware implementation of said algorithm."* (emphasis added)

9.13. The guidelines continue:

"If the mathematical method does not serve a technical purpose and the claimed technical implementation does not go beyond a generic technical implementation, the mathematical method does not contribute to the technical character of the invention. In such a case, it is not sufficient that the mathematical method is algorithmically more efficient than prior-art mathematical methods to establish a technical effect..."

However, if it is established that the mathematical method produces a technical effect when it is applied to a field of technology and/or adapted to a specific technical implementation, the computational efficiency of the steps affecting that established technical effect is to be taken into account when assessing inventive step."

9.14. Part G-II 3.3.1 of the EPO Guidelines explores AI patentability specifically and explains that *"artificial intelligence and machine learning are based on computational models and algorithms for classification, clustering, regression and dimensionality reduction, such as neural networks, genetic algorithms, support vector machines... Such computational models and algorithms are per se of an abstract mathematical nature, irrespective of whether they can be "trained" based on training data."*

9.15. The EPO Guidelines also provide examples of what AI specific inventions may or may not be patentable by virtue of their technical contribution:

"For example, the use of a neural network in a heart-monitoring apparatus for the purpose of identifying irregular heartbeats makes a technical contribution. The classification of digital images, videos, audio or speech signals based on low-level features (e.g. edges or pixel attributes for images) are further typical technical applications of classification algorithms. Classifying text documents solely in respect of their textual content is however not regarded to be per se a technical purpose but a linguistic one (T 1358/09). Classifying abstract data records or even "telecommunication network data records" without any indication of a technical use being made of the resulting classification is also not per se a technical purpose, even if the classification algorithm may be considered to have valuable mathematical properties such as robustness (T 1784/06)."

9.16. In a seminar in January 2020, the UKIPO explained that AI inventions typically fall into three categories: 1) applications of AI to particular fields or problems, 2) techniques for training the AI, and 3) the core AI itself. Of these, the UKIPO recognised that inventions in categories 1 and 2 are likely to be patentable, but those in category 3 could be considered more problematic where they are claimed independently of an application that has the required technical character.

Conclusion

9.17. As a practical matter, there seems little reason to avoid applying the principles laid down in *Aerotel*, *Symbian*, *AT&T* and *HTC*, to establish the circumstances in which an AI invention is patentable. The unhappy reality however is, that those principles also extend the problems inherent with software patents to the more specific area of AI.

9.18. The categories of excluded subject matter, therefore, undoubtedly create challenges when seeking to patent AI related inventions.

9.19. While case law and patent office guidelines in particular, have gone to lengths to clarify the limits imposed by those restrictions, there is little ability for patent owners and practitioners to reliably predict the outcome of decisions relating to computer implemented inventions. Given the costs associated with obtaining patent protection, it is obviously desirable for such unpredictability to be removed as a feature of the patent system in the UK. In a digital world, businesses of all types and sizes are innovating using data, machine learning and AI. It is important for the UK economy that potential users of the patent system are not discriminated against because they lack the sophistication, experience or financial

resources needed to understand how excluded subject matter can impact the patentability of their inventions from the outset.

- 9.20. There has been significant commentary on the utility of the computer program exclusion in s.1(2) PA 1977, particularly in light of the Art 27 of TRIPS³¹ which states that *“patents shall be available for any inventions, whether products or process, in all fields of technology, provided that they are new, involve and inventive step and are capable of industrial application.”*
- 9.21. In light of the importance of digital innovation in today’s society, we think it is now difficult to resist the logic of that discussion and in particular the conclusion that it is sufficient to simply rely on the patentability requirements of novelty and inventive step to determine whether a computer program, and so an AI related invention, is patentable and to remove the computer program exclusion altogether.³²
- 9.22. Should computer programs ‘as such’ be removed from the list of excluded categories, this leaves the examination of AI related inventions subject to the same patentability requirements as other inventions and avoids the need to draw any line. Stimulating innovation in AI is likely to be achieved by ensuring that (1) potential inventors understand that patent protection is, in principle, available in relation to AI-related technology and that (2) the means for assessing whether such an innovation is patentable is both affordable and repeatable.
- 10. Do restrictions on the availability of patent rights cause problems for ethical oversight of AI inventions?**
- 10.1. Not directly. It is assumed that this question relates to the need to sufficiently disclose any invention in a patent specification such that a skilled person may be able to put it into practice. Specifically, it appears to ask whether excluded categories of the subject matter, and the related case law, reduces the oversight of AI presumably by making it unattractive for inventors to pursue patent protection and, as a result, avoid the need to disclose their technologies. It is unclear however, the extent to which published applications are in fact used to monitor the general progress of computer-based technologies let alone the ethical application of technologies that are the subject of those applications.
- 10.2. Further, it should not be assumed that the patent system helps to ensure that developments in technology are properly examined with respect to their ethical merit.
- 10.3. The Patents Act 1977 provides that a patent should not be granted for an invention the commercial exploitation of which would be contrary to public policy or morality.³³ The EPO Guidelines apply a test “to deny protection to inventions likely to induce riot or public disorder, or to lead to criminal [or offensive] behaviour.” However, as highlighted by the Technical Board of Appeal in T356/93 making plants and seeds resistant to weeds and fungal diseases by changing their genome was not held to be in breach of Art. 53(a) EPC 1973, despite it being opposed on the grounds that the exploitation of the invention was likely to cause serious damage to the environment. Such a practice could not be considered to be wrong, as such, in the light of the conventionally accepted standards of conduct of European culture. Plant biotechnology *per se* could not be regarded as being more contrary to public morality than traditional selective breeding.
- 10.4. Similarly, AI’s use as a tool to produce a patentable invention could be said to be in line with accepted cultural standards around, for example the processing of data. The fact that a human could use such a data processing system for an immoral purpose is unlikely to lead to a successful challenge under Art 53(a).
- 10.5. Put another way, the removal of these patentability restrictions is unlikely to increase the ethical oversight of AI, apart from the overall fact that it **may** give increased exposure to technical

³¹ World Trade Organisation’s Agreement on Trade-Related Aspects of Intellectual Property Rights

³² Q255-RES-2017-en: AIPPI Resolution on Patentability of Computer Implemented Inventions, 17 October 2017

³³ S. 1(3) based on Art 53(a) EPC’s exclusion relating to “ordre publique or morality”

developments in the field because more applications are being pursued as a result. That increase in patent filings is by no means a certainty, however, in that examining patentability focusses on the technical character of the invention applied for rather than looking to highlight any potential human use cases that may contravene Art 53(a) and patents are normally drafted with that in mind.

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?

11.1. It may pose a challenge in the future, but only in respect of certain types of AI and even in those circumstances, there may be ways to overcome such a challenge.

11.2. As it currently stands and from the data at the start of this paper, hopeful patentees certainly feel they have sufficient enabling disclosure to support the claims of their AI-related patent applications. In terms of new filings, the second fastest-growing field at the EPO in 2019 was computer technology (+10.2%), the driving factor for which was the increase in patent applications related to artificial intelligence, especially in the areas of machine learning and pattern recognition, image data processing and generation, and data retrieval.³⁴

Legal Principles

11.3. The principles of enabling disclosure and sufficiency are critical to the effective operation of any patent system. In addition to being a central tenet of the patent bargain, under UK law an enabling disclosure is needed to provide 'support' to the patent claims (s.5(2)(a) PA), it is a requirement of a valid application (s.14(3)) and it is essential to one of the grounds of revocation (s.72).³⁵

11.4. Kitchen J (as he was then) provided a summary of general principles relating to the requirement of sufficiency in *Eli Lilly v Human Genome Sciences*³⁶:

- i. The first step is to identify the invention and that is done by reading and construing the claims;
- ii. In the case of a product claim that means making or otherwise obtaining the product;
- iii. In the case of a process claim that means working the process;
- iv. Sufficiency of the disclosure must be assessed on the basis of the specification as a whole including the description and the claims;
- v. The disclosure is aimed at the skilled person who may use his common general knowledge to supplement the information contained in the specification;
- vi. The specification must be sufficient to allow the invention to be performed over the whole scope of the claim³⁷;
- vii. The specification must be sufficient to allow the invention to be so performed without undue burden

11.5. Aldous J *usefully* summarised the nature of the person skilled in the art in *Mentor Corp v Hollister Inc*³⁸.

"The section requires the skilled man to be able to perform the invention. Such a man is the ordinary addressee of the patent. He must be assumed to be possessed of the common general knowledge in the art and the necessary and expertise to apply that knowledge. He is the man of average and intelligence, but is not expected to be able to exercise any invention. In some arts he may have a degree, in others he will be a man with practical experience only. Further, in circumstances where the art encompasses more than one technology, the notional skilled

³⁴ <https://www.epo.org/news-events/news/2020/20200312.html>

³⁵ Lord Hoffman, *Biogen v Medeva* [1997] RPC 1

³⁶ [2002] RPC 41

³⁷ As contrasted with *Biogen v Medeva* where the patent adequately enables a limited subset of embodiments but its claims are broader.

³⁸ [1991] FSR 557

person will be possessed of those technologies which may mean that he will have the knowledge of more than one person."

- 11.6. In addressing the extent of disclosure necessary to satisfy the requirement of sufficiency, in *Halliburton v Smith*³⁹, Pumfrey J usefully commented:

"The sufficiency of a specification is a question of fact and necessarily depends upon the nature of the invention and the attributes of the skilled person. There is no general rule, and although statements like... the skilled person must be enabled to perform the invention without prolonged research, enquiry and experiment' give a flavour of the problem they do not really help... This is a particular risk where the subject of the specification is very complex and its development would anyway be expected to be accompanied by a great amount of work. What is 'prolonged' in this context? It is always necessary to keep a balance between the interests of the public and the interests of the patentee in the sense that it is necessary to guard against imposing too high a standard of disclosure merely because the subject matter is inherently complex."

- 11.7. What would be considered 'prolonged research' (or 'undue burden' in *Eli Lilly*) was explored by the Court of Appeal in *Halliburton* which stated that "the setting of a gigantic project, even if merely routine, will not do."⁴⁰

- 11.8. Sufficiency was considered in detail by the Supreme Court in the recent decision of *Regeneron v Kymab*⁴¹ which focussed in particular on product claims, the technical contribution to the art and the requirement that a claim that contained a range must be enabled across the entirety of that range. Specifically, Lord Briggs provided for 8 steps, namely:

- i. the extent of the monopoly conferred by the patent should correspond with the extent of the contribution which it makes to the art. (Art 83 EPC);
- ii. in the case of a product claim, the contribution to the art is the ability of the skilled person to make the product itself, rather than (if different) the invention;
- iii. patentees may choose how to frame the range of products for which they claim protection but must ensure that they make no broader claim than is enabled by their disclosure;
- iv. the disclosure required of the patentee coupled with the common general knowledge existing as at the priority date, must be sufficient to enable the skilled person to make substantially all the types or embodiments of products within the scope of the claim;
- v. a claim which seeks to protect products which cannot be made by the skilled person using the disclosure in the patent will, subject to de minimis or wholly irrelevant exceptions, be bound to exceed the contribution to the art made by the patent, measured as it must be at the priority date;
- vi. rather than demonstrating that every embodiment within the scope of the claim has been tried, tested and proved to have been enabled to be made, patentees may rely upon a principle of general application if it would appear reasonably likely to enable the whole range of products within the scope of the claim to be made. But they take the risk, if challenged, that the supposed general principle will be proved at trial not in fact to enable a significant, relevant, part of the claimed range to be made, as at the priority date;
- vii. nor will a claim which in substance passes the sufficiency test be defeated by dividing the product claim into a range denominated by some wholly irrelevant factor, such as the length of a mouse's tail. The requirement to show enablement across the whole scope of the claim applies only across a relevant range. Put broadly, the range will be relevant if it is denominated by reference to a variable which significantly affects the value or utility of the product in achieving the purpose for which it is to be made; and

³⁹ [2006] RPC 2

⁴⁰ [2006] EWCA Civ 1715 at paras 17 and 18.

⁴¹ [2020] UKSC 27

- viii. Enablement across the scope of a product claim is not established merely by showing that all products within the relevant range will, if and when they can be made, deliver the same general benefit intended to be generated by the invention, regardless how valuable and ground-breaking that invention may prove to be.

Providing an enabling disclosure in respect of AI-related inventions

- 11.9. On the face of it, technical experts do not believe there is any particular challenge in communicating a new and inventive AI technology in a way that allows a skilled person to make the product in question or work the process, without the requirement of prolonged research and without undue burden.

“While some AI-related inventions do not easily conform to a typical reductionist approach of explainability, and though aspects of AI technology can introduce randomisation, it is by no means impossible to have an effective exposition of an AI -related invention in a document such as a patent and for any such invention to be repeatable. Indeed it may be possible to tell people how to build a particular AI system on the back of a postcard, but the challenge is normally that it requires significant resources to execute.... [t]he algorithms and techniques employed by many of the newer high-profile AI systems have been around for a while but it is the application that is new and, in particular, the increase in sophistication that is created by the sheer scale of the data used. GPT-3 is a good example of this.” [Professor David Barber – Director of the Centre for AI, University College London, November 2020]

- 11.10. This would seem to be supported by statistics relating to AI applications and grants. There seems little problem in granting inventions relating to AI with the just the top 10 AI patent owners having over 4,880 enforceable patent rights in Europe alone.



	Germany	United Kingdom	France	Ireland	Italy	Switzerland	Belgium	Spain	Austria	Denmark	Next 30	TOTAL
Google LLC	155	92	74	57	48	50	49	43	38	44	274	924
Siemens AG	77	53	52	19	51	31	20	27	29	20	334	713
Samsung Electronics Co Ltd	49	56	44	25	18	19	19	16	16	16	391	669
Intel Corporation	75	103	59	40	26	28	26	19	19	20	215	630
Nokia Corp	39	32	32	13	13	11	12	13	10	11	216	402
Accenture Ltd.	20	20	20	20	9	14	11	10	7	7	209	347
Toyota Motor Corp	84	44	49	16	17	15	13	13	11	11	41	314
Huawei Technologies Co., Ltd.	51	50	46	21	23	16	15	19	13	14	44	312
Mitsubishi Electric Corp	57	33	29	19	16	17	17	16	13	13	62	292
Sony Corp	48	42	40	13	11	14	14	12	12	10	62	278
Next 9990	432	303	263	114	108	104	101	88	72	71	896	2,552
TOTAL	1,087	828	708	357	340	319	297	276	240	237	2,744	7,433

Number of patent families that have at least one current grant in the respective country, by organisation.

- 11.11. As Pumfrey J highlighted in *Halliburton*, much depends on the nature of the invention. If the AI-related invention is complex, then it might necessarily come with significant technical support in the description of the specification. But this is not a problem specific to AI and one that today’s patent system is equipped to handle, as highlighted by the table above.

- 11.12. The question of undue burden is likely to be an important one when assessing the nature of disclosure in AI-related patents. Specifically, must the skilled person undertake a ‘gigantic project, even if routine’ in order to arrive at the AI-related patented invention? It must be expected that any elements of randomisation or minor inconsistencies either in training data or introduced by small variations in implementation will only hinder those seeking to re-work the claimed invention.

- 11.13. That said, much will depend on specific wording of the claim and the skilled person's own experience and ability to determine those elements which are important to follow and those which will not impact practising the claimed invention. Indeed, it may be that with the use of AI tools of the near future, such an invention may be easily understood and implemented by the nominal skilled person.
- 11.14. Much like challenges in patenting computer implemented inventions, success will come down to appropriate patent drafting and ensuring that it is properly enabled across the full scope of the claims.

Depositing

- 11.15. To the extent that the inherent nature of AI prevents an enabling disclosure for example, due to randomness or unavoidable minor deviations or that it is simply impossible to say how the invention works because of some black box effect, all may not be lost. The problem of a patent specification being unable to provide an enabling disclosure due to the very nature of the subject matter is not a new problem.
- 11.16. Section 125A PA, titled "Disclosure of an invention by specification: availability of samples of biological material", implements the Budapest Treaty⁴² and provides for deposits of biological material at an international depositary authority in order to meet the legal requirement of sufficiency of disclosure. The use of rare micro-organisms in the 1940s and 1950s to create new antibiotics gave rise to problems of enabling disclosure because such organisms were not generally available to the public. In 1970, the US Court of Customs and Patent Appeals⁴³ confirmed the need for a deposit requirement because "inventors could not sufficiently disclose how to obtain the [new microorganism] from nature, [t]he screening process was unpredictable [and as] a result, a great deal of time and experimentation might be necessary... If so, this would constitute "undue experimentation" [and] the enablement requirement is not satisfied."
- 11.17. Notwithstanding that DNA sequencing is now more reliable and that advances in technology that have improved predictability in the development of new micro-organisms, the international depositary authorities established by the Budapest Treaty remain very much in use today with some 82 countries being a signatory. At the very least, the Budapest Treaty provided an important solution to a sufficiency problem at a time when technology could not be used to help describe inventions in a repeatable manner.
- 11.18. The question of whether a similar system of disclosure is appropriate for AI-related inventions would need careful consideration. The cost and practicality of creating such a depositing system need not be as onerous as that required for biological materials, however uniquely the underlying AI that is deposited is likely to be subject to copyright protection which causes its own complications. For example, how would the owner provide a licence to third parties to use the deposited software and on what terms? If the licence is very broad (which out of necessity it may be), an owner may be reluctant to effectively forgo its copyright protection for the sake of patent protection. What happens when the patent expires – does the copyright licence to the deposited AI continue? It may be simpler and cheaper for an owner of an AI simply to keep the black box functionality as a trade secret and instead commercialise the output rather than the item itself.

AI as a tool

- 11.19. Even amongst technical experts such as Professor Barber, AI is still very much perceived in terms of being a tool used by humans, rather than having significance as its own independent origin of creativity. On the issue of whether the use of the AI need be disclosed at all in a patent application (for example in relation to an innovation in a non-AI related sphere), there seems little need to stray far from existing practices relating to the use of tools in developing new technologies. Whether they need to be disclosed

⁴² Budapest Treaty on the International Recognition of the Deposit of Microorganisms for the Purposes of Patent Procedure, 1977

⁴³ In re Argoudelis 434 F.2d 1390.

in patent specifications depends on the nature of the invention, the nature of the support needed for the claims and what the skilled person, imputed with CGK, would understand was required in order to practise the invention. Does the mention of a mass spectrometer in a patent specification require the exact calibration parameters to be disclosed or should the skilled person be taken to have that knowledge, at least insofar as it does not affect their ability to practise the invention? To not mention a mass spectrometer at all in the early days of mass spectrometers may have risked an objection relating to sufficiency but it is unlikely to do so now. Similar considerations will likely have to be taken into account for any AI that is used as a tool.

Open or closed

11.20. Although it is not confined to matters of sufficiency, an open source software model might equally apply to AI related IP in the same way it applies to copyright in computer software. In such a case, the underlying AI and potentially the training data can typically be published on a platform such as GitHub for download and used under a form of open source licence. It has been shown that patentability and earning royalty returns on proprietary (copyright and patent) software are not the not only ways for owners to extract value from their digital innovations (e.g. Open Core, Support, SaaS). As alluded to earlier, with acquisitions such as IBM's purchase of RedHat, there is growing evidence that computer-based innovation can support both a proprietary and open source model, and there is no reason to believe AI should be any different.

11.21. However, open source regimes are not free from their own problems. It is less easy to derive value when compared to a proprietary licensing model and it is not uncommon for innovations to be published and heralded as publicly available or open source only for implementers to find that on a detailed inspection, critical elements are missing.

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?

12.1. Within the UK's legal framework, there are already mechanisms in place to provide for the disclosure of certain information, which may include the details relating to an AI invention. Such disclosure obligations may be subject to qualifications, for example for reasons of national security⁴⁴, but as the use and complexity of AI grows so may the need to balance the public interest⁴⁵ and provide more detailed transparency and explainability of how AI has been used to arrive at a particular outcome.

12.2. In its report "AI in the UK: ready, willing and able?"⁴⁶, the House of Lords Select Committee on Artificial Intelligence considered the importance of 'intelligibility' which is described as referring both to the ability of experts to understand the technical workings of AI (technical transparency) and the ability for the AI to demonstrate the reasons for taking a particular course of action (explainability). The paper concluded:

"Companies and organisations need to improve the intelligibility of their AI systems. Without this, regulators may need to step in and prohibit the use of opaque technology in significant and sensitive areas of life and society. To ensure that our use of AI does not inadvertently prejudice the treatment of particular groups in society, we call for the Government to incentivise the development of new approaches to the auditing of datasets used in AI, and to encourage greater diversity in the training and recruitment of AI specialists."

"We believe it is not acceptable to deploy any artificial intelligence system which could have a substantial impact on an individual's life, unless it can generate a full and satisfactory explanation for the decisions it will take. In cases such as deep neural networks, where it is not

⁴⁴ Section 24 Freedom of Information Act 2000

⁴⁵ See s.2 Freedom of Information Act 2000

⁴⁶ HL Paper 100, 18 April 2018

yet possible to generate thorough explanations for the decisions that are made, this may mean delaying their deployment for particular uses until alternative solutions are found.”

- 12.3. The report concluded, inter alia, that “[b]lanket AI-specific regulation, at this stage, would be inappropriate. We believe that existing sector-specific regulators are best placed to consider the impact on their sectors of any subsequent regulation which may be needed.” As is reflected in the House of Lords’ report, data protection regulation is of central importance in the endeavour of ensuring companies and institutions can provide the correct level of accessibility, specifically in relation to personal data. The Information Commissioner’s Office (the UK’s information rights regulator) has looked in detail at the impact of using AI on data protection regulations and has released several reports in conjunction with the Alan Turing Institute – the UK’s national institute for data science and artificial intelligence.

Data Protection

- 12.4. In June 2019, the ICO and the Alan Turing Institute published its report on Project Explain⁴⁷ which addressed practical guidance for organisations, to assist them with explaining AI decisions to the individuals affected. The report identifies three key themes that emerged from the research:
- the importance of context in explaining AI decisions;
 - the need for education and awareness around AI; and
 - the various challenges to providing explanations.

- 12.5. The ICO’s AI-related analyses run to many publications and a summary of their content is outside the scope of this answer. However, insofar as it is relevant to the disclosure of sufficient detail of an AI system the ICO, in response to its own consultation in February 2020⁴⁸, stated:

“The Data Protection Act 2018 (DPA 18) gives us the power to carry out various auditing and investigation activities. We believe that, in some cases, this includes the recovery and analysis of evidence, including the AI systems themselves.”

Oversight by industry regulators

- 12.6. In line with the recommendations made by the House of Lords Select Committee, individual regulators have embarked upon a review of the impact of AI on their own industries and focussed on the need for transparency and explainability. Below we look at efforts being made in relation to regulators in financial services and healthcare and medicines.
- 12.7. In respect of financial service institutions, as a result of a survey conducted by the Bank of England and the Financial Conduct Authority⁴⁹, the FCA is now working with the Alan Turing Institute⁵⁰ in respect of addressing transparency standards in AI, with a report due in early 2021.
- 12.8. Given the focus on transparency and auditability of decision making in financial services, the problem is partly addressed by existing financial services obligations relating to risk management. In particular model risk management normally forms part of a bank’s enterprise risk management framework and provides a useful basis on which to assess the adequacy of AI of ML systems.⁵¹
- 12.9. Medicine and healthcare is another industry where the transparency and explainability of AI is a critical issue given the potential for serious harm. In 2019, the Medicines and Healthcare Regulatory Agency (MHRA) released an AI position paper with the BSI and the Association for the Advancement of Medical

⁴⁷ <https://ico.org.uk/media/2615039/project-explain-20190603.pdf>

⁴⁸ <https://ico.org.uk/media/about-the-ico/consultation-responses/2618057/ai-guidance-consultation-responses-20200730.pdf>

⁴⁹ <https://www.fca.org.uk/publication/research/research-note-on-machine-learning-in-uk-financial-services.pdf>

⁵⁰ <https://www.fca.org.uk/insight/ai-transparency-financial-services-why-what-who-and-when>

⁵¹ <https://www.pwc.co.uk/data-analytics/documents/model-risk-management-of-ai-machine-learning-systems.pdf>

Instrumentation.⁵² The paper which focussed on developing new AI related standards, framed the issue of explainability and disclosure in terms of product validation. Specifically:

“Validation of AI solutions is an area where new approaches will be needed to address the challenges that these technologies bring, when compared to traditional software. Two principle reasons for this have been identified.

1. *AI developers will need to have access to appropriate datasets to train their solution during the design and development phases. It is important that these datasets contain adequate quantities and variety of information, and that they are appropriate for the intended use of the solution.*
2. *There is a need for effective supervision of continuous learning systems. Active learning presents a challenge because an algorithm can change outputs in response to receiving new data over time. Unintended or erroneous responses to new data inputs could be more difficult to identify under these conditions. Version control (or similar oversight procedures) would need to be put in place to effectively manage these risks.*

The [relevant] standard specifies requirements for the procedures and recording of product validation. Different approaches to validation for AI solutions are likely to be developed, depending on the proportionate levels of risk. These approaches could include analytical and clinical validation methods. How an output is reached by an AI solution must be clearly explained to ensure professional confidence and trust.”

Legal Proceedings

- 12.10. It is further foreseeable that in order to establish whether a party has committed some unlawful act, either in civil or criminal legal proceedings, it is necessary to understand more fully the role of any AI system that has been involved. To achieve that level of understanding it may be necessary for the entire AI system to be disclosed, and there is nothing in principle that prevents such disclosure under the UK courts existing rules. However, even such a blanket disclosure may not be determinative. As Arm highlighted in the House of Lords Select Committee report:

“what happens when a genuine AI machine makes a decision which results in harm? In such cases unravelling the machine’s thought processes may not be straightforward”.

- 12.11. The purpose of disclosure will presumably be to help determine the allocation of liability in such proceedings, a topic which is further explored below.

National Security and Investment Bill

- 12.12. On 11 November 2020, the UK government published its plans for a National Security and Investment Bill⁵³. The proposed legislation would give the UK government certain powers to intervene in foreign direct investment in the UK, in a manner similar to the CFIUS regime in the United States.⁵⁴ The intention is for the government to take a “targeted, proportionate approach to ensure it can scrutinise, impose conditions on or, as a last resort, block a deal in any sector where there is an unacceptable risk to national security.” Under the framework, investors and businesses will have to tell the government about proposed deals in a limited number of sensitive sectors, such as defence and AI, and its screening powers will be extended to include assets and intellectual property as well as companies. AI is specifically called out as a technology sector in respect of which a UK investment will require mandatory notification.
- 12.13. Section 19 of the bill provides a broad power for the Secretary of State to require any information in relation to the exercise of the Secretary of State’s functions under the legislation. Therefore if there is

⁵² <https://www.bsigroup.com/globalassets/localfiles/en-gb/about-bsi/nsb/innovation/mhra-ai-paper-2019.pdf>

⁵³ <https://www.gov.uk/government/news/new-powers-to-protect-uk-from-malicious-investment-and-strengthen-economic-resilience>

⁵⁴ <https://home.treasury.gov/policy-issues/international/the-committee-on-foreign-investment-in-the-united-states-cfius>

some reason to suspect that the functioning of AI may cause a national security concern, then any investment may trigger a right for the government to ask for details of its operation, which may, in principle, extend to a disclosure of its underlying functionality.

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?

13.1. We think it unlikely the level of inventive step will change in such a way that it cannot be accommodated by current patent law. Using AI as a tool may allow a skilled person to access the state of the art in a more efficient way, but when compared with existing tools such as internet search engines, the issue appears to be a matter of degree than the principles currently embodied in UK law.

13.2. Inventive step is present if the invention “is not obvious to a person skilled in the art, having regard to any matter which forms part of the state of the art”.⁵⁵

13.3. In reformulating the famous test in *Windsurfing*, Jacob LJ⁵⁶ set down the following steps for assessing inventive step:

(1)(a) Identify the notional “person skilled in the art”

(1)(b) Identify the relevant common general knowledge of that person;

(2) Identify the inventive concept of the claim in question or if that cannot readily be done, construe it;

(3) Identify what, if any, differences exist between the matter cited as forming part of the “state of the art” and the inventive concept of the claim or the claim as construed;

(4) Viewed without any knowledge of the alleged invention as claimed, do those differences constitute steps which would have been obvious to the person skilled in the art or do they require any degree of invention?

13.4. To the extent that an innovation relating to AI is new, there is no reason to assume that, *prima facie*, the current legal framework is not sufficiently capable of making a determination of inventive step. There are, however, areas that may be of significant relevance to the issue of AI related inventions.

Technical obviousness

13.5. Patents are generally not available for commercial improvements⁵⁷ but rather technical innovation over the state of the art. As a matter of policy, the public should be entitled to practise that which has been made available to the public, with obvious modifications.⁵⁸

13.6. This may be an important distinction where AI algorithms may be trained or adapted in a new field, simply because, until that particular AI existed, it was not deemed to be commercially useful to do so. In the absence of something more than obvious technical modifications, such a new commercial use should not be enough to satisfy the inventive step requirement.

Skilled Person and Common General Knowledge

13.7. The starting point for assessing who is the person skilled in the art is neatly explained by Laddie J in *Eli Lilly*:

“... obviousness has to be assessed through the eyes of the skilled but non-inventive man in the art. This is not a real person. He is a legal creation. He is supposed to offer an objective

⁵⁵ section 3 Patents Act 1977

⁵⁶ *Pozzoli SPA v BDMO SA* [2007] EWCA Civ 588,

⁵⁷ *Dyson v Hoover* [2002] RPC 22

⁵⁸ *Hallen v Brabantia* [1989] RPC 307

test of whether a particular development can be protected by a patent. He is deemed to have looked at and read publicly available documents and to know of public uses in the prior art. He understands all languages and dialects. He never misses the obvious nor stumbles on the inventive. He has no private idiosyncratic preferences or dislikes. He never thinks laterally. He differs from all real people in one or more of these characteristics.”

- 13.8. In many ways this description is not unlike that of an AI, in that it has access to large number of documents, in any number of languages and has no idiosyncratic preferences or dislikes. That said, the advent of the internet and automated translation tools has exponentially increased the accessibility of documents and information in a similar regard. Accounts have been written as to whether the person of ordinary skill in the art should be replaced by a ‘machine of ordinary skill in the art’ in relation to AI inventions and that it will dramatically affect the level of creativity that the person skilled in the art is expected to possess, to cause the non-obviousness threshold for patentability to rise significantly.⁵⁹
- 13.9. Again, comparisons with the use of the internet to access information are useful. Has this caused such a dramatic rise in the level of inventive step required for patentability? If so that trend is not reflected in the numbers of granted patents being issued by the EPO with more than twice the number of patents being granted in 2019 compared with 2010.⁶⁰
- 13.10. In reality, it is possible that as tools develop and improve the ability to draw on information that could be used to show that an innovation should be obvious over prior art, they may be equally used to help develop and support new and inventive concepts, and reduce the amount of effort needed to decide whether a line of research is worthwhile. As Kitchin J (as he was then) pointed out in *Generics v Lundbeck*⁶¹:

“The question of obviousness must be considered on the facts of each case. The court must consider the weight to be attached to any particular factor in light of all the relevant circumstances. These may include such matters as the motive to find a solution to the problem the patent address, the number and extent of the possible avenues of research, the effort involved in pursuing them and the expectation of success.”

- 13.11. As a result, the inventive step required may indeed vary from case to case and there is no reason to believe that a sophisticated tool, such as an AI system, should be excluded from that analysis if it is something that the notional skilled person would have used.
- 13.12. As regards AI Generated Inventions, there is still a question mark over whether AI has reached level of sophistication to truly create inventions that would have been patentable if created by a human counterpart. A thorough review of AI in this context was published in June 2020⁶² which concluded that many legal commentators over-state the nature of what is possible with today’s technology and that:

“[a] more plausible scenario is, perhaps, where AI is applied as computational methods in the course of solving problems in various fields of research and development. In such situations, however, it appears unclear what degree of AI involvement should be considered to be prejudicial for recognising a human as an inventor, especially, given that the use of problem-solving tools and methods has not been a material factor from an inventorship perspective. (Otherwise, we should also be concerned about situations where microorganisms are used in research and development of biotechnological inventions, as they appear to be more viable candidates to act as ‘autonomous agents’ having consciousness of their own).”

- 13.13. Moreover, technical experts such as Professor Barber consider that the step change in sophistication often has less to do with innovative system architectures and more to do with the sheer volume of data

⁵⁹ ‘From PHOSITA to the MOSITA’, Max Planck Institute for Innovation and Competition, 24 June 2020

⁶⁰ <https://www.epo.org/about-us/annual-reports-statistics/statistics.html>

⁶¹ [2007] RPC 32

⁶² Daria Kim, ‘AI-Generated Inventions’: Time to Get the Record Straight?, GRUR International, Volume 69, Issue 5, May 2020, Pages 443–456, <https://doi.org/10.1093/grurint/ikaa061>

that is used in the training and operation of the AI itself. Professor Barber's example of GPT-3 highlights the point.

- 13.14. A natural language processor, GPT-3 is an example of narrow AI, that is AI that is specific to a task or use and distinct from Artificial General Intelligence - hypothetical intelligence of a machine that can understand or learn any intellectual task that a human being can, believed to be still many decades away⁶³. GPT-3 employs a form of machine learning called 'unsupervised' as the training data does not include an indication of "right" and "wrong", as with supervised learning. Instead it calculates the probability that its output satisfactorily answers the user's question by examining the training data. During this process it assigns "weight" to the algorithm process that provided the correct answer and so records or 'learns' that applying the same method in the future may also yield a correct answer. While this "weighting" process has been around for many years, its ground-breaking performance in GPT-3 is said to be down to the number of weights held and used by the algorithm to process each query, specifically 175 billion – ten times more than any other at time of writing.
- 13.15. As to whether this AI system itself would be patentable, it is unclear what aspects of GPT-3 would be inventive when compared with previous state of the art systems, particularly if it is highly likely, and so obvious, that increasing the data processing capabilities in terms of weights would also increase the accuracy of its output.
- 13.16. What may have been patentable are aspects of the architecture that relate to the weighting mechanism itself (subject to the usual patentability requirements). However, it remains to be seen whether an inventor would have known that such an innovation was worth patenting without understanding how it would operate with large data sets. Indeed, without that understanding would such a patent filing run the risk of being insufficient (following *Regeneron*) because at the filing date, it would not have been possible to have made the product in the way it now currently operates?
- 13.17. As regards AI Generated Inventions, GPT-3, as a very current example of a sophisticated AI, still seems far from creating inventions that would in theory be possible of patent protection.
- 13.18. More likely is that GPT-3 could be used to process large amounts of data in order for the output to be useful in a broader inventive context (i.e. AI as a tool).

Mosaicing

- 13.19. Combining different prior art references to show that an invention is obvious is only permissible to the extent that a skilled person, confronted with a particular citation would turn to some other citation to supplement the information provided by the first.⁶⁴
- 13.20. An interesting question does arise as to whether the notional skilled person has an AI tool such as GPT-3 at their disposal. Where the data source is Common Crawl's petabytes of internet based information, and it is easily accessible through a GPT-3 type interface, it may become easier for experimenters to look to parallel technologies in a broader range of industries than was previously possible within the same constraints.
- 13.21. Again however, this appears to be an issue already confronted by the advent of the internet's increased access to information, rather than a new point of principle. The impact of AI on issues such as this and for example, on changing standards around what would be 'obvious to try'⁶⁵ for the skilled person in the context of assessing inventive step, are ultimately a question of fact. There is little evidence that the current law on obviousness is poorly equipped to deal with AI related inventions, at least as it relates to AI as a tool.

⁶³ PE634.420: "How Artificial Intelligence Works", European Parliamentary Research Service, STOA, March 2019.

⁶⁴ *Glaxo Group Ltd's patent* [2004] RPC 43

⁶⁵ *Actavis UK Ltd v Merck & Co Ltd* [2008] EWCA Civ 444

13.22. When AI Generated Inventions do become a reality, the question of what the invention actually is and how it is claimed comes into sharp focus as do issues of sufficiency. Particularly if AI has generated a new molecule or product, but with little understanding of why or how (because of black box opacity, for example) it may be difficult to judge obviousness or understand what was the CGK should be in the specific circumstance – e.g. because it is unclear whether the AI has utilised material considered to be outside the skilled person's CGK in order to invent. However, in challenging such an inventive step, it is open to parties to demonstrate that using such an AI to generate the invention would have been obvious at the priority date and its output would render the claim in question obvious.

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

14.1. No, for the reasons outlined above in our answer to question 13.

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

15.1. Patent infringement is a tortious act with strict liability⁶⁶. Under UK law a *person* infringes a patent by doing those things specified in section 60(1) PA, without the consent of the proprietor. Further others may also infringe, for example joint tortfeasors or as two parties having a common design⁶⁷.

15.2. As a non-human, AI lacks the necessary legal personhood to be held liable for patent infringement, much like any other offence.

15.3. The only party that may be so liable is that who has control of the AI itself. For AI to create something more than just information output, it is likely that there will need to be some form of human intervention that would give rise to infringement. Take for example, an AI that could generate a design of a mechanical device in response to an engineering problem and was set up so that its solution was automatically downloaded to a 3D printer, and the resulting product was an infringement of a pre-existing patent. Who has made the 3D print out? It is clearly arguable that the human is infringing, as the AI was set up as a tool at the disposal of the human and directed by the human to print, albeit automatically.

15.4. The human as the infringer may argue that they are not liable to pay damages as an innocent infringer under s.62 PA, with success depending on whether they can establish that they were not aware that the patent existed and also they had no reasonable grounds for supposing that such patent existed. This is a scenario that exists today and does not particularly turn on any issues introduced by AI.

15.5. In any event, in most circumstances loss will be incurred due to the commercialisation of the infringing product or process. It is difficult to see how such sales activity is done without human direction and so it is always likely to provide a suitable person against which a patent may be enforced.

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

16.1. In theory yes, but such problems may not be insurmountable. Should the AI be suspected of undertaking a patent process in a way that is obscured from disclosure through its own internal black box working, there may be problems evidencing such operation in infringement proceedings. However, complexity in this area is not new. Court procedures relevant the conduct of patent litigation in the UK provide for disclosure mechanisms and the submission of experiments in order to help establish whether infringement as matter of fact exists. Further, it may be that ongoing work into the explainability

⁶⁶ For a general discussion of AI and civil liability see paragraph 10.4 to 10.6 in our response to the Designs questions.

⁶⁷ Mustill LJ in *Unilever v Gillette* [1989] RPC 583

of AI helps provide a framework to determine the nature of black box operation in a more transparent way.

- 16.2. On one level, it could be argued that the very operation of AI is nothing more than complex algorithmic calculations and so should any patent be said to be solely infringed by such operation then it may be subject to a validity 'squeeze' in that it must also be excluded subject matter under s.1(2)(a) PA.
- 16.3. The bigger problem may be where the operation of the AI is combined with some other functionality which, when taken together, amount to patent infringement.

COPYRIGHT AND RELATED RIGHTS

Questions

The use of copyright works and data by AI systems

1. **Do you agree with the above description of how AI may use copyright works and databases, when infringement takes place and which exceptions apply? Are there other technical and legal aspects that need to be considered?**
 - 1.1. We broadly agree with the consultation description of how AI may use copyright works and databases, when infringement takes place and which exceptions may apply.
 - 1.2. One aspect which requires some further consideration is the impact of questions of copyright ownership on the rights of assignees, licensees and others who are incorporating works generated by or using AI technologies into their own original works. This is particularly important in creative industries such as the audio-visual and music sectors where it is becoming increasingly common that works produced using AI technologies are used in the creation of other (or parts of other) original creative works.
 - 1.3. The film and television industry is a good example of a sector where we often see works generated using AI technologies incorporated into other creative works. AI technologies can significantly contribute to the creation of audio-visual works, notably through technology that assists with script writing, music, CGI and special effects. The producer of a film or television programme will acquire or license works created by a number of third parties, including works created using these AI technologies, to use in the film or television programme. Producers therefore rely on clear rules regarding ownership, assignment, and licensing of copyright works not only to ensure that the audio-visual works they produce will not give rise to claims of copyright infringement, but also to attract financing for production.
 - 1.4. Questions of ownership can be tricky when it comes to content created using AI technologies as there can be multiple parties involved. AI technologies can be complex and there is the potential of multiple contributions from designers, programmers, operators and owners, who may all have some influence in the output of an AI system that creates a single work. However, UK copyright law is well equipped to deal with complexity and contains pragmatic mechanisms for providing certainty to users of works in tricky situations (e.g. employers are designated as the first owners of copyright of works created during the course of employment⁶⁸). The general respect for contractual freedom in the UK is another key aspect in this regard.
 - 1.5. The CDPA provides that ownership of a computer-generated work (i.e. a work which has no human author) belongs to the legal person who undertook the arrangements necessary for its creation⁶⁹. There is significant debate over whether this provision covers works generated by AI (see response to question 5). It is important that, whichever approach is ultimately taken, the rules regarding assignment and licensing remain clear for those who lawfully acquire the rights to incorporate such material into their own creative content. Clarity of ownership is of key importance in the creative industries so that prospective assignees, licensees, rightsholders or users can be certain that they have acquired the necessary rights for their intended use of the work.
 - 1.6. That being said, whilst AI technology is advancing, it is still in its infancy. In our experience AI technologies in the context of the content sector are primarily being used as tools to help creators to achieve their own creative goals rather than the AI itself creating the entire work. As noted in the

⁶⁸ CDPA, section 11(3)

⁶⁹ CDPA, sections 9(3) and 178

consultation description, where AI is simply used as a tool and the output of that tool can be traced back to human creativity, the work produced will generally be protected in the same way as any other work.

- 1.7. The copyright system in the UK generally functions well and has been successful in accommodating numerous technological changes over the past few decades. Considering the current state of AI technology and the uncertainty over how it may evolve over time, restraint is warranted when considering whether changes to the law are currently necessary.

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

- 2.1. Based on what we currently know about AI technology, there is no immediate need to make changes to UK law and there is no reason to believe that existing liability doctrines are inadequate to deal with new scenarios engendered by the use of AI technology. UK copyright law is generally fit for purpose and the use of existing liability principles to address legal issues that may arise in connection with the current use of AI technologies is appropriate.
- 2.2. It is very important to ensure that the use of AI technology does not become a way for a party to escape liability simply because an AI system is involved. The existing doctrines of direct, secondary and accessory liability (when applied in line with their purpose of holding people responsible for infringing acts they commit or cause others to commit) provide clear and adequate guidance for allocating liability in the AI context.
- 2.3. Section 16(1) of the CDPA sets out the acts restricted by copyright in a work. These acts include reproducing a copyright work, which includes the making of temporary copies by electronic means. This is particularly relevant in the context of AI technology as such technologies may create temporary copies of works as part of the input process which “trains” the AI. There are a number of permitted reproductions of copyright works, but these are unlikely to apply where AI is used in a commercial context.
- 2.4. Section 16(2) CDPA suggests that AI cannot itself infringe the copyright in a work, as copyright can only be infringed by a natural or legal person who does, or authorises another to do, one of the acts restricted by copyright, most likely the designers, owners, and operators of the AI. This suggests that authorisation is the natural route to liability for infringement by AI. As AI cannot technically infringe copyright itself, the creator or operator of the AI will be liable instead.
- 2.5. The current approach functions well as it provides the owner of the copyright with a natural or legal person against whom they can bring infringement proceedings if necessary. Designers, owners and operators of AI systems who do one of the acts restricted under section 16(1) may be subject to liability for primary infringement of copyright and/or secondary or accessory liability if others use their systems to directly infringe copyright. An example of the court applying the liability regime in the context of online platforms is the *Newzbin* case⁷⁰ where the High Court found that Newzbin, an online indexing and search service, was liable for infringements of copyright committed by its users who had used its service to help them download copyright-protected films and other media. The Court held that Newzbin authorised the copying of those films by its users, was a joint tortfeasor with such users (having procured and engaged with its users in a common design to copy those films) and was primarily liable for communicating those films to the public.⁷¹

⁷⁰ *Twentieth Century Fox and others v Newzbin Limited* [2010] EWHC 608

⁷¹ Newzbin is still the leading authority on this matter and has been applied in the following cases: *Dramatico Entertainment Ltd v British Sky Broadcasting Ltd* [2012] EWHC 268 (Ch); [2012] E.C.D.R. 14; *EMI Records Ltd v British Sky Broadcasting Ltd* [2013] EWHC 379 (Ch); [2013] E.C.D.R. 8; *Paramount Home Entertainment International Ltd v British Sky Broadcasting Ltd* [2013] EWHC 3479 (Ch); [2014] E.C.D.R. 7.; and *Warner Music UK Ltd & Sony Music Entertainment UK Ltd v TuneIn Inc* [2019] EWHC 2923 (Ch); [2020] E.C.D.R. 8.

- 2.6. In our view, the existing liability doctrines under UK law establish a general, well-accepted framework for analysing claims of primary and secondary/accessory copyright infringement in the context of new technologies.
3. **Is there a need to clarify existing exceptions, to create new ones, or to promote licensing, in order to support the use of copyright works by AI systems? Please provide any evidence to justify this.**
- 3.1. In terms of existing exceptions, the consultation refers to the concept of fair dealing and the text and data mining exception (TDM).
- 3.2. Both exceptions are fairly limited. The TDM⁷² only permits the making of copies for the purposes of text and data mining for non-commercial research and requires the party using the AI system to have lawful access to the material in question. Whilst limited, this exception strikes a balance between allowing AI technologies to grow and develop for important research purposes and protecting rightsholders by ruling out a number of AI systems which may attempt to use the copyright works for commercial gain without providing fair remuneration to the rightsholder. In our view, any move towards a broader TDM exception, as has been introduced in the EU by Article 4⁷³ of the DSM Copyright Directive (DSMCD), requires an abundance of caution. The potential damage to rightsholders must be considered along with the potential application of international copyright rules against formalities that must be complied with as a pre-condition to the enjoyment of copyright protection.
- 3.3. The other fair dealing exceptions may be relevant in specific situations, including for the purposes of research or private study, criticism, review, quotation, news reporting and caricature, parody or pastiche⁷⁴. Whether a fair dealing exception applies will always be a question of fact, degree and interpretation.
- 3.4. AI systems using copyright works for commercial purposes will likely have to rely on licensing. It could be argued that this may inhibit innovation in AI technology by providing obstacles to the research and development of AI. However, it is not clear that the possibility of copyright infringement by AI systems is a major impediment to their use as shown by the huge advances in AI technology up to this point.
- 3.5. The core copyright principles adopted by the UK have balanced the rights and responsibilities of copyright owners and users through decades of technological change. The UK must be wary of further extending copyright exceptions and thereby tilting the balance against the creation, production and protection of new copyright works.
4. **Is there a need to provide additional protection for copyright or database owners whose works are used by AI systems? Please provide any evidence to justify this.**
- 4.1. As noted above, the copyright system in the UK generally functions well and has been successful in accommodating numerous technological changes over the past few decades. The principles adopted by the UK have generally sought to balance the rights of copyright owners and users.
- 4.2. Case law shows that the courts will intervene where an imbalance is found. The High Court overturned new copyright legislation that would have allowed users to create personal copies of copyright-protected materials without permission from rightsholders⁷⁵.

⁷² CDPA, section 29A

⁷³ Article 4 of the DSMCD is a general TDM exception which is not limited to non-commercial use, although rightsholders are permitted to 'opt out' of the exception by reserving their rights through machine-readable or other means. This effectively means rightsholders would have to 'opt in' to get full copyright protection.

⁷⁴ CDPA, sections 29, 29A, 30 and 30A

⁷⁵ See *R (on the application of British Academy of Songwriters, Composers and Authors) v Secretary of State for Business, Innovation and Skills* [2015] EWHC 1723 (Admin); [2015] EWHC 2041 (Admin) where the UK government sought to implement art 5(2)(b) of the 2001 Copyright in Information Society Directive by introducing regulations (the Copyright and Rights in

- 4.3. In addition, most commercial copyright content is protected by technological measures and the CDPA provides a number of civil and criminal remedies in respect of acts which circumvent the protection which has been applied to a copyright works other than a computer program by way of “effective technological measures”. In particular, there are civil remedies available where effective technological measures have been applied to a copyright work (other than a computer program) and a person circumvents those measures. For these purposes, “technological measures” are defined as: *“...any technology, device or component which is designed, in the normal course of its operation, to protect a copyright work other than a computer program.”*⁷⁶
- 4.4. AI technologies are also being used to administer rights. These technologies can assist rightsholders with the licensing process to help ensure they receive remuneration for the use of their work. In this way AI technologies can also be used to protect rightsholders.
- 4.5. It does not appear that there is a current need to change the law to provide additional protection for copyright or database owners whose works are used by AI systems but this could change if new exceptions were to be introduced or the existing exceptions were broadened.
5. **Should content generated by AI be eligible for protection by copyright or related rights?**
- 5.1. As noted above, regardless of how ownership principles develop for works produced with the use of AI technology, it is important that the rules regarding assignment and licensing remain clear for those who lawfully acquire the rights to incorporate such material into their own creative content.
- 5.2. Where AI is simply used as a tool and the output of that tool can be traced back to a human the work produced will generally be protected in the same way as any other work. Copyright will protect the work to the extent that it is the human creator’s “own intellectual creation⁷⁷” and the first owner of the work will be the creator.
- 5.3. As discussed, the CDPA provides that ownership of a computer-generated work (i.e. a work which has no human author) belongs to the legal person who undertook the arrangements necessary for its creation⁷⁸. Protection lasts for 50 years from the end of the year in which the work is made⁷⁹. The rationale behind this provision is to create an exception to the requirement of human authorship and to provide recognition for the work that goes into creating an AI technology capable of generating works independently. A broad interpretation of this provision is preferred, where works generated by AI technology are protected by copyright, however, we acknowledge that there is much debate in this area.
- 5.4. The key issue is whether the originality requirement, that the work must be the creator’s “own intellectual creation”, excludes works that are generated solely by AI because AI does not have legal personality. However, if one of the key rationales behind copyright law is to encourage creativity by providing a reward to creators, it might seem illogical that the humans who have created AI capable of generating works independently would receive no reward for their creativity in doing so. On the other hand, across the world it is clear that non-human authors, including animals, are not permitted to be authors of copyright works⁸⁰ and it should be pointed out that the human that created the AI generating creative works still receives reward through their ownership of the copyright in the AI software itself.

Performances (Personal Copies for Private Use) Regulations) in August 2014 which inserted a new provision (s 28B) into the Copyright, Designs and Patents Act 1988. However, the absence of a balance to all parties affected - particularly rightsholders - led to judicial review proceedings being successfully brought to challenge the legality of the implementing regulations.

⁷⁶ CDPA, section 296ZA and 296ZF(1)

⁷⁷ Or was produced independently by the expenditure of sufficient skill, labour, knowledge etc. under the traditional English law approach

⁷⁸ CDPA, sections 9(3) and 178

⁷⁹ CDPA, section 12(7)

⁸⁰ The “Monkey selfie case” ([Naruto v. Slater, No. 16-15469 \(9th Cir. 2018\)](#)) in the US determined that there was no copyright in pictures taken by a monkey because the pictures were taken without any human intervention. Works created by animals are not entitled to registration under the US copyright system.

- 5.5. These are difficult questions, however, UK copyright law, unlike many other copyright systems, has put in place a provision to deal with this. A broad interpretation of section 9(3) CDPA would solve many problems by making it clear that AI generated works are subject to copyright and the owner of these works is the legal person who made the necessary arrangements for the AI creation. Additionally, if the creators of AI technologies can be liable for copyright infringement it seems unfair that they would not also be rewarded through copyright for their own creations of original works. Not providing a sufficient reward through copyright might disincentivise investment in AI technologies as it may be difficult to protect and monetise the work generated by the AI. That being said, there are systems other than copyright which can provide protection to creators of AI, such as patent law.
- 5.6. Ultimately, based on the current state of AI technology, there will be a legal person somewhere who is responsible for the machine and the existing copyright framework can deal with this. Therefore, restraint is warranted when considering whether changes to the law are necessary at this point in time. In any case, however ownership of AI generated works develops, the rules regarding assignment and licensing must remain clear.
6. **If so, what form should this protection take, who should benefit from it, and how long should it last?**
 - 6.1. As noted above, if a broad interpretation of section 9(3) CDPA is accepted then these questions are already answered within the existing legal framework.
 - 6.2. As discussed above, questions of ownership can be tricky because of the collaborative nature of AI systems as there will likely be multiple parties involved in creating and running an AI system. In straightforward cases, it will likely be the organisation that owns and operates the AI system that will own the output of the AI system. The existing ownership principles of copyright law can provide guidance here and on a practical level, there will generally be contracts between the parties involved which will provide the clarity necessary as to who owns the output of the AI system.
7. **Do other issues need to be considered in relation to content produced by AI systems?**
 - 7.1. No reply.
8. **Does copyright provide adequate protection for software which implements AI?**
 - 8.1. As noted in the consultation description, software is protected as a literary work under section 3(1)(b) CDPA and it is generally accepted that most AI software will be protected by copyright. It appears that the primary legal protection for software is copyright. However, as highlighted in our answer to the patent questions in this consultation, we are aware that there are an increasing amount of AI-related patent filings in the UK, which may provide an alternative or additional method of protection for certain AI-related inventions.
 - 8.2. Similar issues to those discussed in question 5 arise where software that implements AI can be said to have itself been created by AI or where the software evolves or is trained by the AI it implements to the extent that it can no longer be said to be derived from an existing copyright work or the result of arrangements undertaken by a legal person. In our experience this is not a huge problem at present. Considering the current state of AI technology an argument can generally be made that any software that has been created, modified or improved by AI is either a modification or improvement of an existing copyright work or it has evolved from arrangements taken by a legal person. The usual principles of copyright would then apply.
 - 8.3. Therefore, we reiterate that restraint is warranted when considering whether changes to the law are necessary at this point in time.

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

- 9.1. In our experience, copyright and copyright licensing does not currently create any unreasonable obstacles to the use or development of AI software. As noted in the consultation description, many AI tools and interfaces are made available under open licences in light of the aim to encourage the development and growth of AI technologies. However, as software is generally protected by copyright, businesses have the freedom to license AI tools however they wish. This provides adequate protection for software creators by allowing them to benefit financially from their work if they wish to.

DESIGNS

Questions

We anticipate changes would be required to legislation if a decision were made to allow AI to be recognised as the author of a design.

1. Do you agree with the analysis above which concludes that it is not possible for AI to be the author or owner of a UK or Community design?

- 1.1. Yes, we agree that it is not currently possible for AI to be the author/designer or owner of a UK or Community design. Whilst the legislation which governs the identity of the author/designer and the owner of a UK or Community design does not explicitly prohibit the recognition of AI in this way, the question is a matter of interpretation of the relevant instrument. In our view, in order to be recognised as the author/designer or owner of a design, AI would have to possess legal personality (which it does not),.

UK registered designs

- 1.2. For UK registered designs, it is important to keep the concepts of “author” and “owner” separate. The author of a design is the “person who creates it” (s.2(3) RDA). However, where a design is generated by computer “in circumstances such that there is no human author, the person by whom the arrangements necessary for the creation of the design are made shall be taken to be the author (s.2(4) RDA). As is the case in relation to UK unregistered designs (see below), the RDA appears to contemplate that the author of the designer will be an individual or individuals, rather than a legal entity.
- 1.3. The author shall be treated as the original proprietor of the design (s.2(1) RDA), except where a design is created by an employee in the course of his employment, when his employer will be treated as the original proprietor (s.2(1B) RDA). These statutory provisions clearly contemplate that natural persons (such as where the author is also the original proprietor) and legal persons (such as the design is created by an employee and the relevant employer is a company or partnership) may be the original proprietor of the design.

UK unregistered designs

- 1.4. For UK unregistered design rights, the CDPA refers to the “designer” of a design, rather than an “author”. The designer of a design is “the person who creates it” (s.214(1) CDPA), and in the case of computer-generated designs, “the person by whom the arrangements necessary for the creation of the design are undertaken shall be taken to be the designer” (s.214(2) CDPA). “Computer-generated” is defined as a design that is generated by computer in circumstances such that there is no human designer.
- 1.5. Russell-Clarke & Howe⁸¹ states that the CDPA seems to contemplate that the designer will be an individual or individuals, rather than a legal entity, but points out that if that were the case, it would be usual for a statute to use the word “individual” rather than “person”. They further comment that in the case of computer-generated designs, it is possible that no particular individual or individuals could be identified as the designer, in which case a company may itself be regarded as the designer (paragraph 4-112). The reference in Russell-Clarke & Howe to a company being regarded as the designer appears to be reference to a company whose employees collectively made the arrangements necessary for the creation of a computer-generated design, rather than a reference to a company being regarded as a designer due to its ownership of technology which creates the design autonomously.

⁸¹ *Russell-Clarke & Howe on Industrial Designs*, 9th Ed.

- 1.6. The designer is the first owner of any design right in a design which is not created in the course of employment (s.215(1) CDPA). Where a design is created in the course of employment, the employer will be the first owner of any design right in the design (s.215(2) CDPA). Where a design qualifies for design right protection by reference to the place in which it was first marketed, the person by whom the articles in question are marketed is the first owner of the design right (s.215(4) CDPA). As is the case in relation to UK registered designs (set out above), these statutory provisions clearly contemplate that natural persons and legal persons may be the first owner of the design.

Community registered and unregistered designs

- 1.7. For Community registered and unregistered designs, no distinction is made between the “author” or “designer” of a design and the owner. The right to the design initially vests in the designer (Art 14(1) CDR), except where a design is developed by an employee in the execution of his duties or following instructions given by his employer, when the design shall vest in the employer (Art 14(3) CDR). The term “designer” is not defined in the CDR, but both Arts 14(2) in the context of joint designers and 14(3) in the context of employees, refer to the “development of designs” which suggests that the designer is the person or persons who develop a design. The European Commission’s Green Paper on the Legal Protection of Industrial Design dated June 1991 (Green Paper), which set out many of the principles that were subsequently adopted in the CDR, stated that designs must be the result of “human activity” (paragraph 5.6.1).
- 1.8. The CDR does not contain any provisions that address the ownership of computer-generated designs, but the question was considered by the European Commission in the Green Paper. The Green Paper states that “the requirement that a design be the result of a human activity covers this type of designs and in the same time gives an answer to the question of the entitlement to the right on such designs”. The Green Paper suggests that the generation of a design by computer is one method of design which should entitle the person using the computer to obtain protection. However, as the Green paper was published in June 1991, it is sensible to assume that its authors did not contemplate the possibility that AI might be capable of developing designs absent any “human activity”.
- 1.9. In the light of the above, the position in relation to Community registered and unregistered designs is broadly the same as the position in relation to UK registered and unregistered designs. The CDR appears to contemplate that designers will be an individual or individuals, rather than a legal entity, and that designs may be owned either by either individuals or by legal entities.

2. **Are there, or could there be, any tensions with the current legislation when seeking to register a design or be recognised as the owner of an AI-created design? Who would be the legal entity applying for the rights?**

The RDA and CDPA set out that the person who made preparations for the creation of a computer-generated design is the author or creator (s2(4) RDA and s214 CDPA). If read broadly this provision could give the creator of the AI system rights to be identified as the author of a design, even if they have not inputted data required for it to operate.

- 2.1. Yes, there could be tensions with the current legislation when seeking to register a design or be recognised as the owner of an AI created design.

The current position

- 2.2. For UK registered designs, the person entitled to seek to register the design is the proprietor. Before the Intellectual Property Act 2014 came into force, s.3(2) RDA provided that an application for registration of a UK registered design should be made by the person claiming to be the proprietor of the

design. Russell-Cooke & Howe⁸² suggests that, whilst s.3(2) RDA was repealed by the IPA, it was not replaced, and the law therefore remains the same (paragraph 2-159).

- 2.3. For UK unregistered designs, the person who is entitled to be recognised as the owner of the design is the designer (s.215(1) CDP), unless it was created in the course of employment or the design qualifies for protection by reference to the place in which it was first marketed.
- 2.4. For Community registered and unregistered designs, the right to seek to register the design initially vests in the designer (Art 14(1) CDR), except where a design is developed by an employee in the execution of his duties or following instructions given by his employer, when the design shall vest in the employer (Art 14(3) CDR).

Designs created by AI

- 2.5. It is necessary to distinguish between two types designs created by AI. The first type of AI created designs concerns the situation whereby AI is used to assist in the process of creating or developing a design ("AI Assisted Designs"). The second type of AI created designs concern the situation whereby the AI creates or develops the design autonomously ("AI Generated Designs").
- 2.6. In our view, in relation to UK registered and unregistered designs, the provisions in the existing legislation relating to computer-generated designs adequately address AI Assisted Designs. In that situation, there are two possibilities, depending on whether a "human author" or a "human designer" can be identified:
 - if a human author or designer can be identified, that person will be the author or designer and will be entitled to apply to register or to be recognised as the owner of the design; or
 - if a human author or designer cannot be identified, the person who arranged for the use of the AI to create the design will be the author or designer and will be entitled to apply to register or to be recognised as the owner of the design.
- 2.7. In relation AI Generated Designs, neither of the above situations apply, as there is no human author or designer, and neither did any person or persons make the arrangements necessary for the creations of the design. Further, as the AI is neither a natural person, nor does it have legal personality, the AI itself could not be the author or designer of the design, and therefore could not apply for a design or seek to be recognised as the owner. In effect, the design would be authorless and ownerless.

Creator of the AI system

- 2.8. The question suggests that the provisions in the existing legislation relating to computer-generated designs could be read broadly enough to give the creator of the AI system rights to be identified as the author of a design, even if they have not inputted the data required for it to operate.
- 2.9. In order to read the existing legislation in the way suggested, we assume that the suggestion is that the creator of the AI system in question, even if they have not inputted data required for it to operate, should be taken to be "the person by whom the arrangements necessary for the creation of the design are undertaken shall be taken to be the author/designer". If that is correct, then that person would be entitled to seek to register a design or be recognised as the owner of a design.
- 2.10. However, this result would mean that the person who created the AI system in question would own all of the designs that it creates, and the person that owns the AI system with a view to the AI creating designs autonomously, would not obtain any rights in those designs. If that were the case, it would represent a strong disincentive for designers to use AI systems to create designs autonomously, as

⁸² Russell-Clarke & Howe on Industrial Designs, 9th Ed.

they would not have ownership of, and would therefore not be able to exploit, any designs that were created by the AI.

3. Who should be recognised as the author of a design created by AI where the system has been bought from a supplier, and the buyer has provided input or data to the system? Does the wording of legislation need to be changed?

3.1. Our response to this question is restricted to consideration of designs that have been created or developed autonomously by AI.

3.2. The traditional justification for the grant of design rights is that they reward authors and designers who come up with new designs by granting a monopoly for a defined period of time in which they can exploit and monetise their designs, thereby incentivising further creativity and innovation in design. This aim of design law needs to be borne in mind when considering the normative question of who should be recognised as the author and therefore the first owner of designs created autonomously by AI. Should the UKIPO wish as a matter of policy to ensure that AI Generated Designs may in fact subsist, then it could consider the alternatives set out below

3.3. There are a number of candidates who could potentially be recognised as the author of a design where the AI system has been bought from a supplier, and the buyer has provided input or data to the system, each of which will be considered below:

1. the AI itself;
2. the creator of the AI; and/or
3. the buyer/user of the AI system who has provided input or data to the system.

3.4. For completeness, we also consider the position of the buyer of the AI who has not provided any input or data to the system.

The AI itself

3.5. For designs developed autonomously by AI, the candidate that most closely meets the definition in current legislation of the author or designer of a design, being “the person who creates” it, is the AI itself. However, AI, even AI which is capable of developing designs autonomously, does not do so due to the incentives on offer, nor does it have any intention or capability of exploiting those designs. On the other hand, the designs created by AI may still demonstrate creativity and innovation, and add in a material way to the totality of human knowledge. Accordingly, whilst designs created autonomously by AI may make a material contribution to human knowledge, recognising AI as an author does not accord with the traditional justification for the grant of designs rights and in fact may stifle creativity and innovative by creating designs but not exploiting them.

3.6. Further, in order for AI to be recognised as the author and/or first owner of a design which it created autonomously, it would have to be capable of owning property, and of transferring property to third parties. In order to do either of these things it would have to have legal personality. Consideration of whether or not to grant legal personality to AI is not a question of IP law and, accordingly, we do not attempt to address it here though the challenge this presents is significant.

The creator of the AI

3.7. As set out in relation to question 2 above, if the creator of the AI which is subsequently used by a third party, whether or not that third party has provided input or data to the system, was recognised as the author and/or first owner of all of the designs that it creates, that would operate as disincentive for designers to use AI systems which were created by others to create designs autonomously, as they

would not have ownership of, and would therefore not be able to exploit, any designs that were created by the AI. Whilst the creator of the AI would be incentivised to continue to create AI systems, because they do not direct the AI to create designs, the creator is not being incentivised to create designs per se. Accordingly, recognising the creator of the AI as the author of does not accord with the traditional justification for the grant of design rights.

Buyer/user of AI

- 3.8. The final possible candidate to be recognised as the author of designs that are created autonomously by AI is the buyer/user of the AI system, which can be split into two according to whether the buyer/user uses the AI 'off the shelf' or instead provides input or data to the AI in order to guide its outputs. If they were to be recognised as the author and/or first owner of the designs autonomously created by an AI system, the buyer of the AI system, whether or not they provide input or data, would be incentivised to try to create designs which demonstrate creativity and innovation. According to the traditional justification for the grant of design rights, both types of buyer/user of the AI system should be granted design rights.
- 3.9. Arguably, a buyer/user who provides input into or data to the AI deserves to be more richly rewarded than a buyer/user who does not, because of the additional effort put into the creation of the designs. However, the level of effort required to create a non-computer-generated design does not determine whether protection should be granted to those designs, provided that the relevant design meets all of the usual criteria for validity. Accordingly, it is difficult to see why a distinction should be drawn between a buyer/user who provides input into or data to the AI and one that does not.

Conclusion

- 3.10. In order for any of the possible candidates to be recognised as the author and/or first owner of a design that has been developed autonomously by AI, the current UK legislation would need to be changed. In particular, the provision regarding who can be recognised as the author of a computer-generated work would need to be expanded.

4. Do you consider that legislation should be changed to allow AI systems to be recognised as the author of a registered design or designer of an unregistered design?

- 4.1. No, for the reasons set out in our response to question 3 above.

5. If so, how should we assess when AI stops being a tool programmed by a human and becomes an intelligent entity capable of producing its own IP? What proof or evidence would be required?

- 5.1. Not applicable.

As aspects of the future UK designs framework will be based on concepts currently found in the Community Design Regulation, we would like your views on the following:

6. Unlike UK domestic legislation, the CDR has no provisions relating specifically to computer-generated designs. Does this result in legal uncertainty in relation to authorship and ownership of computer-generated designs? Would the same apply to AI-generated designs?

- 6.1. As set out in our response to question 1 above, although the CDR does not explicitly refer to computer-generated designs, but it appears from the Green Paper that the intention was that computer-generated designs would fall within the regime set out in the CDR, provided that there was at least some element of "human activity". The Green Paper also suggests that, where AI is used to assist in the process of creating or developing a design, the identity of the author and owner of a Community registered or unregistered design would be determined as set out in our response to question 1.

- 6.2. However, the Green Paper does not discuss the position in relation to designs which are created or developed autonomously by AI, which is unsurprising given that it was published in 1991. Accordingly, much like the RDA and the CDPA, the CDR does not cater for this category of designs. This means that the degree of legal uncertainty in relation to the CDR is the same as in relation to the RDA and CDPA.

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

- 7.1. We are not aware of any other issues specific to the CDR which you should consider in relation to AI.

Infringement

8. Can the actions of AI infringe a registered or unregistered design? Can AI do the acts set out in law (s7(2) RDA)?

- 8.1. We anticipate that AI systems will increasingly play a role in the creation of designs. We need to consider the implications on design infringement. This section sets out the relevant law. It asks whether the legal tests for deciding whether infringement has occurred can be applied to AI systems.

- 8.2. A design is infringed when a person does anything which is the exclusive right of the owner without the owner's permission (s7A(1) RDA). This includes making, offering, putting on the market, importing, exporting or using a product a design is incorporated or applied in, or stocking a product for any of these reasons (s7(2) RDA and Article 19(1) CDR).

The RDA and CDR both set out that protection relates to a design and any design which does not produce a different overall impression on the informed user. It also sets out that the degree of freedom available to the author when creating the design is taken into account (S7(1) RDA and Article 10 CDR).

- 8.3. The legislation does not clarify what is meant by an "informed user", but the courts in *Samsung v Apple* and *PepsiCo v Grupo Promer* have provided the following guidance:

- a) they are a user of the product in which the design is intended to be incorporated (not a designer, manufacturer, technical expert, or seller);
- b) they are particularly observant;
- c) they have knowledge of the design corpus and the design features normally included in the designs existing in the sector concerned;
- d) they are interested in the products and show a relatively high degree of attention when they use them;
- e) e) they conduct a direct comparison of the designs (unless that is impractical or uncommon due to specific circumstances or certain characteristics of the products)."

- 8.4. The CDPA contains similar provisions for unregistered UK design rights. It sets out that the owner has the exclusive right to reproduce the design for commercial purposes (s226(1) CDPA). This "means copying the design so as to produce articles exactly or substantially to that design" (s226(2) CDPA). A design right is infringed by a person who does anything which is the exclusive right of the owner without permission (s226(3) CDPA). Case law says that the test to decide if a design is substantially similar is objective. And it should consider whether a customer would think the designs are similar.

- 8.5. Further questions arise when considering who should be held liable for infringing a design. It will be for the courts to determine the answer but the use of "a person" in relation to infringement (s7A(1) and (2) RDA, s226 CDPA and Article 19 CDR) suggests only a legal person may infringe a design.

UK registered designs

- 8.6. S.7(1) RDA gives the registered proprietor of a design the exclusive right to use a design and any design which does not produce on the informed user a different overall impression. S.7(2) RDA states that “use” of design includes making, offering, putting on the market, importing, exporting or using a product in which the design is incorporated or to which it is applied, or stocking such a product for any of these purposes.

UK unregistered designs

- 8.7. S.226(1) CDPA gives the owner of a design the exclusive right to reproduce the design for commercial purposes by making articles to that design, or by making a design document recording the design for the purpose of enabling such articles to be made.

Community registered and unregistered designs

- 8.8. Arts 19(1) and (2) CDR give the holder of a registered or unregistered Community design the exclusive right to use it. Such use includes the making, offering, putting on the market, importing, exporting or using a product in which the design is incorporated or to which it is applied, or stock such a product for those purposes.

Infringement by AI

- 8.9. As a practical matter, of the acts that amount to infringement listed above, it may be possible to say that AI, in some automated manner may be capable of:

- offering or putting on the market a product in which a design is incorporated or to which it is applied contrary to s.7(1) RDA;
- making a design document recording a design for the purpose of enabling articles to be made to the design contrary to s.226(1) CDPA; and/or
- offering or putting on the market a product in which a design is incorporated or to which it is applied contrary to Art 19(1) CDR.

- 8.10. As a practical matter AI is unlikely to be able to do any of the other acts listed above, because doing those acts would involve physical interaction with the infringing product which AI is not capable of. All of those acts would require at least some degree of involvement of a natural person.

9. When considering infringement are there, or could there be, any difficulties applying existing legal concepts in the registered designs framework to AI technology? Does AI affect the use of the “informed user” in measuring overall impression?

- 9.1. We understand this question to be limited to registered UK and Community designs.

- 9.2. In relation to UK and Community registered designs, a design is infringed where a third party does any of the acts listed in our response to question 8 above in relation to that design, or any design which does not produce on the informed user a different overall impression (S.7(1) RDA; Art 10(1) CDR). In assessing whether a design produces a different overall impression on the informed user, the degree of freedom of the designer in developing the design shall be taken into consideration (s.7(3) RDA; Art 10(2) CDR).

- 9.3. One of the possible consequences of the use of AI to create designs is that there could be a proliferation of designs which would greatly expand the state of the art, and make the assessment of whether a design is ‘new’ (which is relevant to the question of validity, not infringement), and/or whether a design

creates the same overall impression as another design (which is relevant to both the question of validity and infringement), more difficult to assess.

- 9.4. The assessment of whether a design is ‘new’ is an objective one, and one in relation to which AI may be able to assist tribunals in assessing the degree of similarity between a design and an alleged piece of prior art. This could be particularly valuable if the use of AI to create designs causes a proliferation of designs and a greatly expanded state of the art. We note that the UKIPO has recently launched (in beta) an AI tool called “Trade Mark Pre-apply” which assists applicants by assessing whether there is already a mark on the register which is too similar to the mark the applicant proposes to apply for.
- 9.5. The assessment of whether a particular product infringes a specific design by producing the same overall impression is more complex because, as set out above, this assessment must be undertaken from the point of view of the informed user, taking into consideration the degree of freedom the designer had when developing the design in question. The deemed characteristics and knowledge of the informed user will be unique to each assessment of infringement. Accordingly, it would be more difficult to use AI to assess whether two designs create the same overall impression.
- 9.6. Where a design is created or infringed by AI, we do not see any reason why the notion of the “informed user” would need to change in order to assess whether there was infringement. In order to address the issue identified above of the possible proliferation of designs due to the use of AI to create designs, could be managed through the use of robust case management, such as by limiting the number of designs that should be considered in order to determine the degree of freedom of the designer, or whether validity is in issue, the number of designs which the applicant for invalidity says establish that the design in question is not new.
- 9.7. The informed user is a legal fiction and is deemed to be a user of the product in which the design is intended to be incorporated. As AI is unlikely to use any such products, we cannot currently envisage any situation in which the informed user would be deemed to be an AI system.
10. **If AI can infringe a registered design, who should be liable for the infringement? Should it be the owner, the programmer, the coder, the trainer, the operator, the provider of training data, or some other party?**
 - 10.1. The infringement of a UK registered or an unregistered design is a form of statutory tort, for which there is strict liability, i.e. it is not necessary for a claimant to demonstrate that the defendant intentionally or negligently infringed the relevant right. The liability regime for Community registered and unregistered designs is based on similar principles.
 - 10.2. As mentioned above, the traditional justification for the grant of design rights is that they reward authors and designers who come up with new designs by granting a monopoly for a defined period of time in which they can exploit and monetise their designs, thereby incentivising further creativity and innovation in design. In order to protect the monopoly granted, the registered proprietor or holder of the relevant registered or unregistered design is permitted to bring proceedings for infringement against those that do the acts listed in our response to question 8 above.
 - 10.3. As set out in our response to question 8 above, AI may be capable of doing some of the acts that constitute infringement. As AI does not have legal personality, the registered proprietor or holder of the relevant registered or unregistered design would not be able to sue the AI for infringement. This means that, unless an AI’s acts of infringement can be attributed to a natural or legal person, or AI was afforded legal personality, the registered proprietor or holder of the relevant design would have no redress for those acts of infringement.
 - 10.4. In July 2020, the European Parliament published a study entitled “Artificial Intelligence and Civil Liability” (Study). The Study advocated the adoption of a risk-management based approach to liability caused by AI. Pursuant to that approach, the Study proposes that liability should attach to the party who is best

placed to identify the risk, control, minimise, and manage it, and that the primary aim of any liability system should be the compensation of victims. The Study suggests that adopting a risk-management based approach does not attempt to directly incentivise the development of safe products, but instead allows such incentives to be provided via market force and regulation, such as reputation and product safety rules. Although not specifically mentioned by the Study, IP law broadly, and designs law specifically, could be one such form of regulation.

- 10.5. On 20 October 2020, the European Parliament adopted a resolution setting out recommendations to the European Commission for a civil liability regime for AI (Resolution). The Resolution proposes that the operator of the AI should be liable for all operations of an AI. The Resolution defines “operator” as both the front-end operator, which is defined as “the natural or legal person who exercises a degree of control over a risk connected with the operation and functions of the AI and benefits from its operation”, and the back end operator, which is defined as “the natural or legal person who, on a continuous basis, defines the features of the technology, provides data and essential backend support service and therefore also exercises a degree of control over the risk connected with the operation and functioning of the AI”. The Resolution further defines “exercising control” as any action that influences the operation of the AI and thus the extent to which it exposes third parties to its potential risks”, and would include determining the input, output or results of the AI or changes to the specific functions or processes within the AI. To the extent that the front-end and back-end operators are not the same natural or legal person, the Resolution proposes that all operators should be jointly and severally liable.
- 10.6. Neither the Study nor the Resolution were developed with IP rights in general, or design rights, specifically, in mind. However, identifying the appropriate natural and legal persons to whom liability for infringing acts committed entirely by AI should attach may be assisted by the idea that liability for the actions of AI should attach to those natural or legal persons who are best placed to identify the risk posed by AI, and control, minimise and manage it. However, we would suggest that the adoption of labels, such as “front-end operator” and “back-end operator” as set out in the Resolution, and the labels set out in the question should be avoided. This is because each label will require a definition, which as can be seen by the definitions set out in the Resolution would tend to be complex. The risk with adopting such complex definitions is that they tend to favour form over substance and may give rise to gaps in protection, and any definitions that are tied too closely to the current state of AI could become obsolete quickly.

TRADE MARKS

Questions

1. **If AI technology becomes a primary purchaser of products, what impact could this have on trade mark law?**
 - 1.1. The premise of the question concerns AI becoming a primary purchaser of products (which we shall assume to mean goods or services of any kind) and our response is framed by that assumption. **However, as a matter of law, without legal personality an AI system is unable to enter into contractual relations and so unable to make purchases.** Only if this significant problem was overcome through amending legislation and if AI were sufficiently technically developed to do so, could this have an impact on trade mark law..
 - 1.2. As the consultation notes, concepts relating to trade mark infringement are founded on human interaction with branding and human involvement in the purchasing process. Trade mark law is therefore based on an underlying premise of transactions being conducted by or for humans. The legal framework, and the way the relevant grounds for infringement and other relevant tests are set are rooted in that underlying assumption. If the nature of that transaction was to change, involving AI as a purchaser, and the human as the end user, the framework of trade mark law may need to adapt to reflect that, placing greater emphasis on, for example, post-sale confusion or other concepts relevant to that transaction scenario. Any potential amendments to trade mark legislation in this respect should be subject to greater and in-depth consultation if and as AI evolves to assume the role anticipated.
 - 1.3. Even if AI could be described as the primary purchaser of products in the context outlined, the impact of a brand on the human consumer/customer/end user and that human's perception of it will remain a core part of the policy framework around which trade mark law is based. Those humans will retain their perceptions of preference, quality and origin, and it is assumed may well retain at least influence if not control over settings or direction of the AI's purchasing decisions, perhaps through excluding and positively desiring certain brands for example. Humans will retain choice, and the ability to choose to purchase something else, from somewhere else, perhaps using different AI assistance, or not using AI at all.
 - 1.4. Nevertheless, based on the present legal framework, the key areas that we consider could be impacted by the question posed would be: (1) the nature, role or characteristics of the average consumer; (2) the assessment of similarity of marks; (3) the assessment of similarity of products; and (4) the assessment of whether there was a likelihood of confusion and/or any relevant injury to a trade mark with a reputation (unfair advantage, detriment to the distinctive character or repute of the mark).

(1) Average Consumer
- 1.5. Question 3 below specifically addresses how AI may affect the role of the average consumer.

(2) Similarity of Marks
- 1.6. The assessment of similarity of marks is presently undertaken by reference to concepts of visual, aural and conceptual similarity, bearing in mind the dominant and distinctive components. This perception is necessarily conducted through human eyes at present, and similarity of marks is a necessary condition for the ground of likelihood of confusion and harm to reputation.
- 1.7. If it was ever possible that an AI could be a primary purchaser of products, presumably the nature of the AI making the purchasing decision would play a relevant role in conducting the comparison of the marks in question at the point of purchase. However, the human end user (for most products) would

still be the ultimate target for most advertising and marketing of a brand, which may in turn influence the scope of AI activity under instructions of the human.

- 1.8. On one hand, it is conceivable that the threshold for similarity of marks could be raised or could become a form of two stage test involving a human assessment as part of the test, but factoring in the ability of AI to better identify differences between marks used for products being chosen to be purchased, depending on how its parameters and instructions and learning is set. On the other hand, the 'AI factor' may be better considered as part of all the relevant circumstances of the case to be assessed when assessing if the relevant harm arises, as referred to below.

(3) Similarity of Goods/Services

- 1.9. At present, the assessment of similarity of goods or services is conducted by reference to established principles of case law, including the nature, intended purpose, method of use, complementarity, competition, distribution channels, relevant public and usual origin of the goods/services.
- 1.10. This is a necessary condition for likelihood of confusion, and can be a relevant factor in assessing conditions for harm to reputation (such as the link or unfair advantage or other injury).
- 1.11. Depending how AI evolved, similarity of goods/services may play less of a role at the stage of the transaction. However, again, the impact of similarity of goods/services would still influence the human ultimate decision maker. Again, whilst trade mark law and policy may need to reflect this change in the transaction, the use of AI may be better considered as a relevant factor to consider when assessing if the relevant harm arises, as referred to below.

(4) Assessment of relevant harm:

- 1.12. The assessment of whether there is a likelihood of confusion or other relevant harm to reputation are multi-factorial assessments based on all relevant circumstances. As noted above, the iterative steps of assessing the average consumer, similarity of marks and similarity of goods or services could be adapted to reflect the role of AI as a primary purchaser. However, if AI does become a primary purchaser, it may be more appropriate to consider this as one of the relevant factors to take into account in the global appreciation of whether the relevant harm arises.
- 1.13. In assessing whether relevant harm arises, established case law sets out a number of principles, based around the underlying premise of a human transaction, as referred to above. These include principles such as 'imperfect recollection' of the marks at issue. Some of these principles would not be applicable to an AI, which could in theory possess potentially endless capacity for storing and analysing information. For example, an AI may be able to perfectly recollect and to compare two marks side by side and would not need to rely on imperfect memory to assess them. This might suggest that the global appreciation test should be simplified, since it would be easier for an AI to appreciate all relevant factors, without going through as many different steps. Similar issues apply to establishing a link, and unfair advantage or detriment to distinctive character and repute.
- 1.14. However, as referred to above, the human user /beneficiary of the AI would necessarily remain the ultimate consumer/customer/end user. The brand would still be (although perhaps not entirely) directed at that human, and depending on how the AI in question was operating, the human may well still be influential or have ultimate control over that AI or purchasing direction (if not transaction itself).
- 1.15. Therefore, even though it may be appropriate to reassess or redefine the assessment of relevant harm as AI evolves, it would still remain necessary to assess the matter at least partly through the eyes of the relevant human consumer/customer/end user. Taking the example of how trade mark law adapted to key word advertising, it may be that trade mark law would assess AI enabled transactions through the eyes of the average user of AI purchasing methods.

- 1.16. The assessment of whether the relevant harm arises may as a result refocus on post-sale confusion / pre-sale confusion. If products and services may be automatically ordered by an AI, whether via voice control or via automatic restocking or recommendation programs, such as the Amazon Dash Replenishment program⁸³ or Alibaba's Fashion AI concept store⁸⁴, where there is no interaction between the consumer and the brand, the concepts of "average consumer" and "likelihood of confusion" may therefore adapt over time, to account more for the impact of advertising and the issue of post-sale harm.
 - 1.17. Unless and until AI evolution and human adoption of such AI technology evolves to a stage of essentially removing human choice, influence or care about what products are supplied (and in particular, their preference, quality and origin), the policy underlying trade mark law still needs to be viewed through a human consumer / customer/ end user lens. Until that time, it may be sufficient and appropriate to adapt existing concepts and, if appropriate create new ones, such as adapting the average internet user into the average AI user, and to bear other AI relevant factors in mind, as part of all the 'relevant factors' of the relevant case, when considering whether the relevant injury arises.
- 2. Are there, or could there be, any difficulties with applying the existing legal concepts in trade mark law to AI technology?**
- 2.1. Please see the answer above.
- 3. Does AI affect the concept of the "average consumer" in measuring likelihood of confusion?**
- 3.1. As set out above, it is important to keep in mind the difference between the AI playing a role in the purchasing act, and the human as the ultimate consumer / customer / end user. It is conceivable in the future that an AI may play a greater role as the 'customer' perhaps in software purchase for systems development or protection, for example, but for vast majority of goods / services humans (or companies controlled by humans) would remain the ultimate consumer / customer / end user.
 - 3.2. Trade mark law concepts already allow for different standards of average consumer based on the nature of the relevant goods or services. For example, average consumers of prescription drugs, expensive products, or complicated systems will pay a higher degree of attention than average consumers of everyday consumer items.
 - 3.3. The involvement of AI could be a relevant factor in the assessment of the average consumer if techniques are developed, as with search engine optimisation (SEO) for example, whereby sellers can influence the likelihood of being selected for purchase by an AI, through as yet unknown techniques. If that was the case, it is possible that could play a role in assessment of the average consumer. However, that may be better assessed as part of all the relevant circumstances when assessing whether the relevant injury arose, in a similar way to how consideration of SEO would be a relevant factor in assessing the relevant injury, but would not equate the average consumer to the search engine.
 - 3.4. As noted above, it may be appropriate to adapt the concept of the average internet user into the average AI user, and to bear other AI relevant factors in mind as part of all the 'relevant factors' of the relevant case, when considering whether the relevant injury arises.

⁸³ <https://www.amazon.co.uk/b?ie=UTF8&node=16463234031>

⁸⁴ <https://fashionunited.uk/news/retail/alibaba-presents-shopping-of-the-future-with-fashion-ai-concept-store/2018070630597>

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

4.1. Section 10 TMA:

*(1) A **person** infringes a registered trade mark if he **uses in the course of trade** a sign which is identical with the trade mark in relation to goods or services which are identical with those for which it is registered.*

*(2) A **person** infringes a registered trade mark if he **uses in the course of trade** a sign where because—*

(a) the sign is identical with the trade mark and is used in relation to goods or services similar to those for which the trade mark is registered, or

(b) the sign is similar to the trade mark and is used in relation to goods or services identical with or similar to those for which the trade mark is registered, there exists a likelihood of confusion on the part of the public, which includes the likelihood of association with the trade mark.

Person

4.2. Section 10 refers to a ‘person’ as the putative infringer, which clearly includes legal as well as natural persons. The underlying Directive 2015/2436 (Directive) refers to ‘all third parties’. As a result neither provision applies to an AI because of the lack of legal personality.

Use in the course of trade

4.3. As a general matter, any use that is made in the course of trade involving AI is likely to be at the direction of a human rather than by the AI itself. Whilst AI lacks the necessary legal personality to be able to infringe, it may be worth considering how interactions between the human and the AI may apply to potentially infringing acts.

4.4. In order to infringe, a trade mark must be used in the course of trade –used in the context of commercial activity with a view to economic advantage and not as a private matter. AI technologies are and will increasingly be capable of selecting or communicating a vast range of products to consumers, in return to a request or search with very few keywords, or even without a request if automatically enabled. Whilst the AI may have been sold or licensed as part of an economic activity, some AI technologies may be seen more as operating akin to an intermediary between products and consumers. However, it may not be clear cut, and whether an AI is using a trade mark in the course of trade may depend on the facts of the case.

4.5. There may be two extremes of AI involvement, with shades of grey in between.

4.6. At one end of the spectrum, an AI may be acting on behalf of the consumer, taking the decisions for the consumer, at the consumer’s request, and be independent from the seller. In that context, if the AI selects products, outside of the scope of any involvement or economic benefit from the sale, but operating as an intermediate service for the consumer, it is difficult to see how the AI would be ‘using’ in the course of trade in the presently understood sense. In that context, the AI is perhaps operating more as a consumer than as the seller.

4.7. At the other end of the spectrum, an AI could be operating on behalf of the seller, as part of the seller’s platform or offering, actively selecting or promoting the sale of goods which benefits the seller directly. In that context, it may be that the AI, either itself, at the behest of the seller, or perhaps the seller itself, would be considered to be ‘using’ the mark in the course of trade (by way of analogy, *Cosmetic Warriors v Amazon*, ([2014] EWHC 181 (Ch)).

- 4.8. It is also conceivable that in between those two extremes, there could be decentralised AIs that may be available for use as independent AIs whose services consumers could use, for no apparent benefit to the AI. It is difficult to foresee all potential possibilities.
- 4.9. As with the example of the average internet user, existing principles from trade mark case law may provide guidance for development in the area of 'use in the course of trade'. In the *Google France*⁸⁵ case, the CJEU drew a distinction between the advertiser on one hand, and Google as the internet referencing service provider on the other. The advertiser picked a keyword identical with a trade mark with the purpose of displaying an advertising link to the site on which the advertiser offered his goods or services for sale. The mark which was selected as a keyword was the trigger for the ad display, and therefore, the court ruled that the advertiser 'used' it in the context of commercial activity. By contrast, a referencing service provider in those circumstances allowed its customers (advertisers) to use signs identical to trade marks, without itself using those signs. In that case, the fact that Google had created the technical condition, necessary for the use of a sign and had been paid for that service, did not mean that it had itself 'used' the sign. Depending where in the spectrum of AI involvement the case in question falls, this case law could inform whether and by whom 'use' takes place.

5. Can the actions of AI infringe a trade mark?

- 5.1. As highlighted above, no. The question above discusses the difficulties arising with an AI falling within Section 10 as presently constituted. However, setting those difficulties aside, in the abstract, the actions of an AI could result in activity which might otherwise infringe a trade mark. The spectrum of AI being 'controlled' or intermediary/consumer as noted above may result in and require different treatment under the law.
- 5.2. Taking the seller-controlled or owned AI, based in the *Cosmetic Warriors* case law, it is conceivable that the actions of an AI may result in infringement and in liability for the seller / operator of the AI (directly, or once on notice). For example, if the actions of an AI resulted in selling and supplying of one product for another that had been ordered without drawing attention to that, or resulted in promoting third party products in circumstances where the user could not tell, or only with difficulty, that the products originated from third parties.
- 5.3. In the grey zone of AI operating without apparent control, it is conceivable that an AI, in particular a decentralised or independent AI, could 'operate' a website only pulling together offers for sale of various products for consumers to view and purchase. Those products could include both infringing products available in the jurisdiction, and also infringing products from outside the jurisdiction and which may not otherwise infringe in their 'home' territory but could be infringing in the consumers territory.
- 5.4. In the latter case, the consumer might not otherwise have happened across the infringing products outside its territory but for the AI. If in that context, the consumer purchases the product for personal use and becomes the importer of those goods, it is difficult to see where potential liability may fall: the seller may be based entirely outside the jurisdiction, the AI may have enabled the advertisement / offer of and even effected the actual sale in the jurisdiction, and the consumer may be the ultimate recipient/user. In such circumstances, the AI may well be enabling otherwise infringing acts, but ones which fall in a legal vacuum. One possible means of addressing this may be to adapt Section 10 to also provide rights holders with the ability to obtain an injunction or other technical means to block sales curated by an AI which would otherwise be unlawful. The rationale may be similar to ISP blocking injunctions (such as obtained in *Cartier & Others*), but may need new legislation or further case law developments, if the underlying acts were not themselves infringing due to the AI not fitting within Section 10.

⁸⁵ C-236/08

6. **If AI can cause trade mark infringement, does this change who could be liable? Should it be the owner, the operator, the programmer, the trainer, the provider of training data, or some other party?**
- 6.1. This is a fundamental question when discussing AI and its impact on trade mark law and liability in general, beyond the law of trade marks. It is very important to ensure that the use of AI technology does not become a way for a party to escape liability simply because an AI system is involved. Due to the private acts of consumers not infringing trade marks, existing concepts of direct and secondary liability may need to be adapted or supplemented by legislation to provide remedies, and where appropriate liability, where AIs commit acts which would otherwise infringe.
- 6.2. The existing liability in Section 10, discussed above, is premised on use in the course of trade and therefore amendments to those provisions or new provisions specific to addressing issues which may arise through AIs may be needed.
- 6.3. Again, it may be appropriate to consider a spectrum of AI involvement.
- 6.4. Taking the seller-controlled or owned AI example, it is logical that the seller should be liable, or capable of becoming liable once on notice, building on the analogies in *Cosmetic Warriors v Amazon* and *L'Oreal v Ebay* case law.
- 6.5. However, where a decentralised independent AI is acting in a manner which would otherwise be unlawful for an economic operator, such as curating and presenting offers of infringing products for sale, a new form of liability or subjection to sanctions, such as injunction or takedown procedures, may be necessary to provide brand owners with a means of obtaining a remedy for problematic actions, as noted in the answer above. It may also be appropriate to consider whether, in those circumstances, the programmer/owner/creator of the AI should be subject to some sort of liability or sanction to procure the cessation of certain acts by the AI, or to injunctions against the release of such AIs in the future, depending on the circumstances. Whilst this latter situation may have some relevance to some released AIs, if AIs are created and released anonymously, it may be difficult to enforce in practice.
- 6.6. If such a decentralised AI was making available goods which already infringed in the jurisdiction, liability for those underlying goods would remain (also) with the source(s) in the jurisdiction in any event.

TRADE SECRETS

Questions

1. **Is trade secret protection important for the AI sector? Does the nature of AI technologies and business influence your answer?**
 - 1.1. Yes, the protection of trade secrets is an important part of the ‘package’ of rights available to those operating in the AI sector.
 - 1.2. The potential limitations of other IP rights of relevance to AI, in particular patents and copyright, serve to highlight the potentially significant advantages of trade secrets as a form of protection to those operating in this field.
 - 1.3. The AI sector is highly innovative. However, whilst patents undoubtedly offer a very significant form of IP protection for inventions, there are several reasons why obtaining patents may not necessarily be the optimal strategy for those operating in the AI sector or, indeed, one that is even commercially viable to them.
 - 1.4. As discussed in more detail in our submissions in response to the patent-specific questions posed in this call for views, obtaining patents over AI technologies themselves can be challenging. It is fair to say that obtaining patents over software implemented inventions (which covers virtually all current AI technologies) can be difficult and is often a long and expensive process, with no guarantee of success.⁸⁶
 - 1.5. Accordingly, the decision as to whether to make the necessary disclosures associated with a patent application, which may itself preclude other forms of protection and in particular protection of information as a trade secret, is one that must be considered very carefully by those operating in the AI sector. Whilst the field is maturing rapidly, a very significant proportion of AI development work is still undertaken by nascent businesses that simply do not have the working capital to pursue patent protection, especially for the type of global portfolio that is necessary to properly protect technology with a truly global market. Accordingly, the risk and cost associated with pursuing patents may simply outweigh the potential benefit for many businesses operating in the AI sector.
 - 1.6. Other limitations with patents in the AI space include the fact that development work is very fast-paced and may change direction rapidly and at a speed that it is simply not practical to track with patent filings and may result in problems with patentability such as those highlighted by *Regeneron*⁸⁷. It is also the case that innovations that are most valuable to AI businesses *commercially* may simply be iterative developments of existing technology, or those which improve on or combine existing/known principles or techniques in a manner which would not in and of itself meet the ‘bar’ for patentability.
 - 1.7. When compared to patent protection, trade secret protection requires very little expenditure save for that associated with properly documenting and ensuring the secure storage of information. It is highly unlikely that those operating in the AI sector will not have access to the requisite knowledge or resource to satisfy these requirements. Moreover, trade secret protection will continue to be available for as long as the requirements for protection are maintained – it does not expire in the same way as a patent – and does not require (indeed, it precludes) disclosure of the protected information to the rest of the world. Accordingly, in principle, trade secret protection can be indefinite.
 - 1.8. As with trade secret protection, copyright protection does not involve any attendant prosecution or registration costs in the UK. However, it is rarely the computer code element of an AI technology itself

⁸⁶https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/817610/Artificial_Intelligence_-_A_worldwide_overview_of_AI_patents.pdf

⁸⁷ (see discussion in our answers to the questions concerning patents)

that is the 'core' aspect that requires protection. Rather, it is typically the algorithms and mathematical models underlying the AI system that is the truly 'clever' and most valuable part of the technology which, generally speaking, will not be protected by copyright. Other innovative aspects of AI technologies, such as the idea of implementing AI in a particular field or combining it with another known technologies or systems may simply not amount to subject matter (in English law, a "work") that is capable of protection under copyright.

- 1.9. There is no such limitation to trade secret protection, the subject of which can be any form of information if it has the requisite quality of confidentiality.
- 1.10. For the reasons outlined above, the relatively low cost and potentially broad and long-term protection offered by trade secret protection is very important for the AI sector.

2. Does the nature of AI pose any problems if UK trade secret protection is required? Does UK trade secret law give adequate protection to aspects of AI technology where no other intellectual property rights are available?

- 2.1. We do not consider that the nature of AI *per se* poses problems if trade secret protection is required in the UK.
- 2.2. However, one issue of relevance is that in order for trade secret protection to be available, the party asserting the right must demonstrate that reasonable steps have been taken in the relevant circumstances to keep the relevant information 'secret'. As noted above, AI is generally a software implemented technology and accordingly, in a climate of ever-increasing cyber-security risk, ensuring that the requisite 'reasonable steps' have been taken will pose an increased challenge over time, as for all other digital sectors.
- 2.3. Further, a number of AI systems are based at least in part on algorithms that are said to be open source (e.g. OpenAI's GPT-3). To the extent that is true of a system that is alleged to be a trade secret, such a system may at least in part be unprotectable.

3. What are the advantages and disadvantages of using trade secrets in the AI sector? Could information that is not shared inhibit AI development?

- 3.1. A number of the advantages and disadvantages of trade secret protection in the AI sector have been touched on above.
- 3.2. Generally speaking, the major advantage of trade secret protection over other IP rights such as patents and copyright is the relatively low cost and potentially broad and long-term protection available.
- 3.3. An obvious disadvantage is that trade secret protection under English law does not offer a complete monopoly over a particular technology. In other words, it does not provide protection against a competitor independently arriving at the same technological solution in the same way that a monopoly patent right does. However, this limitation reflects that there is no disclosure required (or indeed permitted) in order to benefit from trade secret protection, which is a key aspect of the patent monopoly 'bargain'. For an entity seeking to protect its proprietary technology (including those in the AI sector), the fact that potentially indefinite protection can be obtained without the need to make any disclosure may well be considered a worthwhile compromise when weighed against obtaining a complete but time-limited monopoly right.
- 3.4. Whether or not keeping information out of the public domain might inhibit AI development raises the same types of questions as have been debated for many years in respect of the monopoly granted by a patent.

- 3.5. It stands to reason that it will be difficult to advance certain technologies if the basis of the 'state of the art' is not accessible as a reference point from which to work for the world at large. However, the availability of trade secret protection may itself act as an incentive to innovation – the fact that developments can be protected over a potentially very long period of time may encourage significant investment in the hope of significant long-term returns. Moreover, the absence of direct access to certain technologies may in and of itself act as a catalyst for the development of alternative solutions which, again, will encourage general innovation.

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

- 4.1. As highlighted in our answer to question 10 of the patent questions, it is not trade secret protection *per se* that causes 'problems' for the ethical oversight of AI inventions.
- 4.2. There is no doubt that certain applications of AI technologies give rise to potentially significant ethical issues. However, in our view the principle of ethical oversight of such technologies and applications⁸⁸ is not completely at odds with trade secret protection.
- 4.3. An important aspect of the ethical oversight of AI technologies lies with the approach of the creators and developers themselves, which generally speaking is unaffected by trade secret considerations.
- 4.4. Whilst the issue becomes potentially more relevant when independent, external oversight is considered, the requirement that information must be 'secret' in order to be protected as a trade secret is not absolute. Rather, what is required is that the information is not 'readily accessible' to those that ordinarily deal with the relevant type of information in question. Accordingly, as with mechanisms that are in place for (for example) the disclosure of confidential information in litigation, we see no reason in principle why trade secret protection could not be maintained whilst at the same time allowances made for access to relevant information for recognised, independent ethical oversight bodies. For example, a direct confidentiality agreement could be put in place in order to confirm that information disclosed for the purposes of independent, ethical review will remain confidential and will only be used for defined oversight purposes.

⁸⁸ See more on which here: <https://www.gov.uk/guidance/understanding-artificial-intelligence-ethics-and-safety>



wiggin

Principal contact:

Calum Smyth

Partner, Intellectual Property

calum.smyth@wiggin.co.uk

www.wiggin.co.uk

