

AI AND INTELLECTUAL PROPERTY: CALL FOR VIEWS

Dr Enrico Bonadio (City, University of London) – Dr Luke McDonagh (London School of Economics) – Dr Plamen Dinev (Goldsmiths, University of London)

PART I: PATENTS

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

The patent system should provide a fair balance between the need to incentivise AI developers/users and allowing the general public to access the invention within a reasonable period of time. It should also ensure that the ownership of both AI-generated inventions and AI systems does not result in unwarranted monopoly control where a select few global entities dominate the field.

2. Can current AI systems devise inventions? Particularly:

a) to what extent is AI a tool for human inventors to use? b) could the AI developer, the user of the AI, or the person who constructs the datasets on which AI is trained, claim inventorship? c) are there situations when a human inventor cannot be identified?

Computer involvement in the inventive process can range from AI acting as a tool to AI generating an invention with a considerable degree of autonomy. There are claims that AI systems have been “autonomously” devising inventions since at least the 1990s (for instance, Stephen Thaler and John Koza both assert that their AI systems have generated various inventions with no human intervention: see <https://artificialinventor.com/dabus>). However, further expert and empirical evidence may be necessary to substantiate these claims.

If the above assertions are indeed correct, it is conceivable that there may be instances where the machine is no longer a mere tool. In such cases, identifying a human as the actual deviser of the invention may be problematic.

Under current law—where machines cannot be designated as inventors—it appears that the AI’s developer, data trainer and/or user may be able to claim inventorship, provided that a human is designated as inventor.

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

Not recognising AI as inventor could potentially encourage dishonesty (i.e. designating a human as inventor where the AI has generated the invention independently) or cause prospective applicants to rely on trade secrets instead. Nevertheless, as the extent to which AI systems are able to devise inventions with no human intervention is disputed, it would be advisable to obtain further expert evidence before introducing any major legislative changes (such as allowing AI to be identified as the sole or joint inventor).

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?

Failing to credit AI as inventor is unlikely to discourage future inventions or have a negative impact on innovation developed using AI. Machines are not motivated by recognition or the prospect of a patent; humans are, and they are still eligible to apply for protection in such cases provided that a human is designated as the inventor.

Keeping inventions confidential could certainly undermine the core rationale of the “patent bargain” where both the inventor and the general public benefit from the disclosure and commercialization of the invention. This may also harm the public and it does nothing to expand the domain.

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

Recognising AI as inventor—subject to the considerations outlined above—could potentially encourage transparency and honesty in patent applications. This argument is particularly valid in cases where the AI’s developer and/or end user have had little involvement.

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

AI does not have any legal rights at present. If machines are to be accepted as inventors in the future, potential candidates for ownership would include the AI’s owner; programmers; users; data providers; and the first person to recognise the significance of the result.

When determining who should be the owner of the patent, a range of policy factors need to be taken into account. On the one hand, if an AI system generates an invention while being operated by a third party, attributing ownership to the end user could have a negative impact on net social welfare (as developers may intentionally restrict access to their AI altogether); instead, attributing ownership to developers may encourage developers to *enable* access. However, a programmer who merely creates a general-purpose claim-generating code does not actually contribute to an important element of the claimed invention. This code merely serves as a tool which others (i.e. end users) could use to make such contributions. A case-by-case approach is therefore likely to be sensible here.

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

Please refer to Questions 3-4 above.

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

As there are separate questions dealing with patentable subject matter, inventive step and disclosure below, we will only consider novelty here. The main concern in terms of novelty is that AI may dramatically expand the prior art universe and make it extremely difficult to satisfy this criterion, especially as inventive machines are able to generate massive volumes of information. On a more practical level, the proliferation of AI-generated claims could also overburden patent offices and hinder the examination process.

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?

AI is not uniquely problematic in this regard and the current law appears to work reasonably well. It is clear that mathematical methods and computer programs are still patentable, provided that they have the requisite technical effect. This would include a wide range of AI systems and AI-generated inventions (e.g. an AI algorithm in a heart-monitoring apparatus which measures irregular heartbeats would not be excluded due to the presence of a technical effect: see https://www.epo.org/law-practice/legal-texts/html/guidelines2018/e/g_ii_3_3_1.htm).

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

N/A

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?

It is often difficult to describe how an AI system actually works, particularly in cases where machine learning creates algorithms that are so complex that even the programmer is unable to understand how or why they bring about a certain result (i.e. ‘black boxes’). Consequently, if e.g. a configuration is not fully known and cannot be reproduced by others, it will be regarded as being insufficiently disclosed—and therefore not patentable.

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?

N/A

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?

AI is likely to become a standard inventive tool in a wide range of fields. As machines have extensive knowledge of the prior art universe, this may necessitate the introduction of stricter procedures, including raising the bar to patentability. The existing law could accommodate this if the test for inventive step is expanded so as to include a skilled person who has access to AI. Yet, such a change would effectively elevate the notional person’s ability to that of a sophisticated machine: thus, the inventive step test may end up including the routine use of AI by skilled persons as an important aspect to be taken into consideration, and applicants could be required to disclose the specific contribution of the machine as to the invention’s conception. A complete substitution of the concept of the “skilled person” with that of a “skilled computer” could make obtaining a patent significantly more difficult, especially as AI becomes progressively more advanced and a greater range of innovative activities appear obvious to such a system.

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

As stated above, there is certainly an argument in favour of amending the existing “person skilled in the art” benchmark to include a person equipped with AI. The potential issue here is that patent offices may have to examine prior art which is comprehensible to machines but not humans (which could be difficult in practice). Said that, completely substituting the existing test with a “machine trained in the art” standard is unlikely to be required in the near future, given the current state of AI.

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

Failing to hold any party accountable might encourage the use of AI systems for infringement purposes. As AI systems have no legal personality, liability is likely to fall on users, data trainers and/or developers. We caution that holding end users liable may be unfair, especially where they are unable to foresee the risk; moreover, users are often legally unsophisticated individuals. Another possibility would be to hold the AI developer or manufacturer accountable – it is already common practice in patent litigation to sue (and condemn) the producer of the infringing goods. This could be a more appropriate approach in the AI context, especially as programmers are arguably in a better position to foresee the risk. They are also more likely to acquire economic value from the AI (e.g. by selling it to end users).

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

Proving infringement is likely to be especially difficult in cases involving black box processing (as described above). However, we are unable to provide estimates on the size of the problem.

PART II: COPYRIGHT

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?

This is an accurate overview. We have outlined other relevant legal issues below.

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

We agree with the description and the case-by-case assessment outlined in the document.

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.

There are various exceptions (e.g. transient copy and text- and data-mining) that are potentially applicable in the AI context. In terms of clarifications, we are particularly interested in the TDM exception under the 2019 Digital Single Market Directive (at the time of writing, however, it appears that the UK Government is unlikely to bring this into UK law; the UK CDPA also provides a text- and data-mining exception, which is limited to the copying of works for the sole purpose of research for a non-commercial purpose). The limitations to the TDM exception under EU law significantly restrict the opportunities for AI *business-oriented* organisations to carry out such activities in the context of machine learning platforms. On the other hand, research centres, heritage institutions and other similar entities that use AI for non-profit purposes will be free to engage in text and data mining activity. It remains to be seen if (in addition to merely mechanical uses of pre-existing copyright works) non-profit organisations can also engage under this exception with respect to expressive and creative AI uses of the mined data; a clarification may therefore be required.

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.

N/A

5. Should content generated by AI be eligible for protection by copyright or related rights?

It may be inevitable that some form of protection is required. However, copyright regimes—including UK and EU law—might not be the best fit. Considering AI-generated output as copyright may require overstressing and rethinking important rules and principles, such as the authorship and originality requirement, which are at the core of copyright systems (e.g. originality under EU law, which remains applicable in the UK, appears to require some degree of human authorship). The traditional copyright system may not be adequate for another reason: the long duration awarded to works of human authorship (this is especially problematic as AI systems can generate huge amounts of creative works and at a much faster rate than any human).

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

We suggest that a sui generis system could strike a fair balance between the need to incentivise the creation of these technologies (through the offer of exclusive rights aimed at maximising the profits out of the final works) and the need to guarantee that human-made creativity continues to thrive despite machines' competition. The benefit of a sui generis regime (as opposed to using the full scope of copyright to protect such works) would be that right holders could be given only a thin scope of protection, allowing them to prevent others from exploiting exact copies of the machine-generated work. It would essentially be a protection against literal copying only. The scope of protection given by the sui generis right should be backed up by a strict fair use/fair dealing doctrine, and its duration would be very short (e.g. three years).

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

N/A

8. Does copyright provide adequate protection for software which implements AI?

Copyright appears to offer appropriate protection for AI software (as a literary work) and we are not aware of any particular problems in this regard.

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

N/A