PHILIPS

World Intellectual Property Organization
34, chemin des Colombettes
CH-1211 Geneva 20
Switzerland

Email: leo.steenbeek@philips.com
Tel: +31 6 10013525

Date: 2020.02.14

Subject: WIPO Conversation on Intellectual
Property (IP) and Artificial Intelligence (AI)
WIPO/IP/AI/2/GE/20/1

Dear Sirs,

Please find attached our reaction to your consultation.

We hope that our response is helpful.

Yours faithfully,
Philips International B.V.
Intellectual Property & Standards

Leo Steenbeek
Principal IP Counsel

PHILIPS

Issue 1 – Inventorship and Ownership

With regard to the issues "1. Inventorship and Ownership" and "2. Patentable Subject Matter and Patentability Guidelines", we need to take as a starting point the four guiding principles for trustworthy AI as published in the document "Guidelines for trustworthy AI", drafted and published by the AI high-level expert group (HLEG) to the European Commission on April 6th, 2019.

The 1st principle of respect for human autonomy is defined as follows: "*The fundamental rights upon which the EU is founded are directed towards ensuring respect for the freedom and autonomy of human beings. Humans interacting with AI systems must be able to keep full and effective self-determination over themselves, and be able to partake in the democratic process. AI systems should not unjustifiably subordinate, coerce, deceive, manipulate, condition or herd humans. Instead, they should be designed to augment, complement and empower human cognitive, social and cultural skills. The allocation of functions between humans and AI systems should follow human-centric design principles and leave meaningful opportunity for human choice. This means securing human oversight over work processes in AI systems. AI systems may also fundamentally change the work sphere. It should support humans in the working environment and aim for the creation of meaningful work.*"

The 2nd principle of prevention of harm is defined as follows: *"AI systems should neither cause nor exacerbate harm or otherwise adversely affect human beings. This entails the protection of human dignity as well as mental and physical integrity. AI systems and the environments in which they operate must be safe and secure."*

The 4th principle of explicability is defined as follows: "*Explicability is crucial for building and maintaining users' trust in AI systems. This means that processes need to be transparent, the capabilities and purpose of AI systems openly communicated, and decisions – to the extent possible – explainable to those directly and indirectly affected. Without such information, a decision cannot be duly contested. An explanation as to why a model has generated a particular output or decision (and what combination of input factors contributed to that) is not always possible. These cases are referred to as 'black box' algorithms and require special attention. In those circumstances, other explicability measures (e.g. traceability, auditability and transparent communication on system capabilities) may be required, provided that the system as a whole respects fundamental rights. The degree to which explicability is needed is highly dependent on the context and the severity of the consequences if that output is erroneous or otherwise inaccurate*".

Also the Expert Group to the European Commission on Liability and New Technologies – New Technologies Formation, states in its key findings that "*It is not necessary to give devices or autonomous systems a legal personality, as the harm these may cause can and should be attributable to existing persons or bodies.*"

PHILIPS

7.(i)   Should the law permit or require that the AI application be named as the inventor or should it be required that a human being be named as the inventor? In the event that a human inventor is required to be named, should the law give indications of the way in which the human inventor should be determined, or should this decision be left to private arrangements, such as corporate policy, with the possibility of judicial review by appeal in accordance with existing laws concerning disputes over inventorship?

Our reply: The implication of the 1st ethical principle on respect for human autonomy is that AI should "*augment, complement and empower human cognitive, social and cultural skills*" and can therefor never be an inventor of a patent. Inventors are human beings. AI augments human skills and expertise but does not make human contribution any less relevant. If an invention is made with the aid of AI, the human who uses AI to find a solution to a problem should be considered to be the inventor. Especially in the complex world of healthcare, AI tools such as machine learning require human domain knowledge and experience to recognize the need and technical problem in need of a solution, put outcomes in context and derive inventive subject matter. In medical research, it is therefore essential to bring together AI engineers, data scientists, and clinical experts. This helps to ensure proper validation and interpretation of AI-generated insights. Furthermore, input and curation of data and insights are critical components of the invention.

In case an invention is being made with the support of an AI system, then the 4th principle of explicability requires that "*processes need to be transparent, the capabilities and purpose of AI systems openly communicated, and decisions – to the extent possible – explainable to those directly and indirectly affected. Without such information, a decision cannot be duly contested.*" In other words, when an invention is being made with the support of an AI system, that the road to come to the invention needs to be transparently explained in the registration process. Only with human oversight, AI can be used, and a human inventor who is responsible for the oversight of the use of the AI application should be named as inventor.

The usual criteria for inventorship therefore apply in relation to inventions supported by AI: whoever contributed to providing a solution to a problem, is an inventor. This raises a question as to which natural persons should be considered to have contributed to the "conception of an AI invention." We believe that the long-standing definition of such conception still applies, e.g. the "formation in the mind of the inventor, of a definite and permanent idea of the complete and operative invention, as it is hereafter to be applied in practice." as defined by the USPTO with reliance on established case law in the US. A couple of scenarios are possible in the context of AI-supported inventions:

- when an invention itself resides in insights or output generated by an AI system in a particular context, whoever recognized an inventive solution to a problem in such insights or output, is an inventor. Whoever engineered or trained the AI system that generated the invention may then be a co-inventor if such work was done within the same context and with the specific aim of obtaining useful output in an effort to solve a recognized problem.
- when an invention focuses on a particular AI system using a generic AI algorithm, whoever conceived the idea of the algorithm and worked out its implementation details sufficiently, is an inventor. Further contributions to the conception of the

PHILIPS

invention may come from recognizing a problem in need of a technical solution, as well as from curation and weighing of training data.

7.(ii)   The inventorship issue also raises the question of who should be recorded as the owner of a patent involving an AI application. Do specific legal provisions need to be introduced to govern the ownership of autonomously generated AI inventions, or should ownership follow from inventorship and any relevant private arrangements, such as corporate policy, concerning attribution of inventorship and ownership?

Our reply: In view of the above, ownership of inventions made by means of an AI is also clear, as the established rules can be followed. We do not see a need for additional regulation/different legal provisions and consider the current law and interpretation sufficient.

7.(iii)   Should the law exclude from the availability of patent protection any invention that has been generated autonomously by an AI application? See also Issue 2, below.

Our reply: No, we feel that this law does not need to change because current guidance on handling computer implemented inventions is sufficient. Most if not all AI related inventions are computer implemented, such that the same guidance can be applied. Article 27(1) of the WTO TRIPs Agreement should be followed: Subject to the provisions of paragraphs 2 and 3, patents shall be available for any inventions, whether products or processes, in all fields of technology, provided that they are new, involve an inventive step and are capable of industrial application. Paragraphs 2 and 3 of Article 27 of the WTO TRIPs Agreement do not allow for an exclusion of inventions made by means of AI.

Issue 2 – patentable subject matter

8.(i)   Should the law exclude from patent eligibility inventions that are autonomously generated by an AI application? See also Issue 1(iii), above.

Our reply: See 7(iii). We do not believe that AI tools should be named as inventors. AI solutions can in principle provide you insights which potentially result in a patentable invention. This would always involve a human inventor that can be identified for each invention.

8.(ii)   Should specific provisions be introduced for inventions assisted by AI or should such inventions be treated in the same way as other computer-assisted inventions?

Our reply: No specific provisions are needed. We prefer to apply the existing patent law. Indeed, the provisions / guidelines / case-law relating to computer-implemented inventions are usually very well applicable to inventions using AI.

PHILIPS

8. (iii)   Do amendments need to be introduced in patent examination guidelines for AI-assisted inventions? If so, please identify which parts or provisions of patent examination guidelines need to be reviewed.

Our reply: No specific provisions are needed. We feel that the current way of working with guidelines on interpretation creates flexibility and is sufficient to address future developments. In the context of the 2nd principle for trustworthy AI, when an invention is being made with the support of an AI system, the road and process to come to the invention need to be transparently explained. We expect that this can be covered adequately by existing provisions requiring the invention to be sufficiently disclosed for it to be carried out by a skilled person. See further our answers to 10.(i) and 10.(ii).

Issue 3 – inventive step or non-obviousness

9. (i)   In the context of AI inventions, what art does the standard refer to? Should the art be the field of technology of the product or service that emerges as the invention from the AI application?

Our reply: Implication of the 1st principle of respect for human autonomy is that AI should "*augment, complement and empower human cognitive, social and cultural skills*". Therefor AI cannot replace the human "*skilled in the art*". The human should be knowledgeable in the respective domain, augmented by AI and be able to trace the process and explain its decision-making process to come to the result which may lead to a registered invention. It is a normal aspect in an inventive step assessment that other fields are considered, if it is obvious in view of the problem to be solved to look into those other fields. For the remainder, again, no new rules are needed. For example, the prior art should continue to be everything that is publicly available prior to the priority date.

9. (ii)   Should the standard of a person skilled in the art be maintained where the invention is autonomously generated by an AI application or should consideration be given to replacing the person by an algorithm trained with data from a designated field of art?

Our reply: There are certain things a skilled person would consider when faced with a certain problem. In the future the skilled person might apply as a standard certain AI tools to improve technology known from prior art. For example, if you merely replace conventional processing techniques by a conventional AI neural network and a skilled person can do this without modifications, e.g. without specific data handling/selection, then this step of AI neural network processing might be considered obvious. The reference is always the skilled human person and his or her standard toolbox at his/her availability to solve the problem underlying the invention. See also our comments as to 9.(i).

Also in respect of AI based inventions we argue that the bar of inventive step should not be lowered but should be maintained at the (perceived high) level as it is today.

PHILIPS

9. (iii)   What implications will having an AI replacing a person skilled in the art have on the determination of the prior art base?

Our reply: As indicated above we are of the opinion that AI will augment a human being and will not replace a person skilled in the art.

9.(iv)   Should AI-generated content qualify as prior art?

Our reply: As far as it is publicly disclosed and enabled, AI generated content should be part of the prior art – in line with the current definition of prior art.

Issue 4: Disclosure

10.(i)   What are the issues that AI-assisted or AI-generated inventions present for the disclosure requirement?

Our reply: Sufficiency of disclosure should continue to be governed by Article 5 PCT, which is incorporated by reference into the PLT as a requirement relating to content. We are of the opinion that no new rules are needed or desired. Having said that, the patent offices should do a thorough review and have expertise to assess the disclosure requirement and enablement of AI assisted or AI generated inventions.

The requirement of Article 5 PCT should be taken seriously, which explains why it is repeated in Article 29(1) WTO TRIPs Agreement.
Article 5 PCT The Description: The description shall disclose the invention in a manner sufficiently clear and complete for the invention to be carried out by a person skilled in the art.

10.(ii)   In the case of machine learning, where the algorithm changes over time with access to data, is the disclosure of the initial algorithm sufficient?

Our reply: As an example, adapting regulators are patentable today, provided there is a good description of the training methodology or how an output of an algorithm will change over time based on input parameters. The claims should be supported in such a way that the skilled person can carry out the invention based on the description.

In that regard, it is important to note that in view of sensitivity and complexity of the AI training methodology, a conventional view of sufficiency of flowchart-based descriptions to enable persons skilled in the art to make and use claimed software inventions, may not be true for inventions in the AI domain. For those inventions, disclosure of, for example, particular neural network models or training data can be necessary to comply with the requirement of the sufficiency of description.

PHILIPS

10.(iii)   Would a system of deposit for algorithms, similar to the deposit of microorganisms, be useful?

Our reply: AI algorithms should not be treated different than regular algorithms for which no deposit system exists. Following our answer under 10.(ii), contrary to inventions involving biological material such as micro-organisms, it seems generally feasible to completely describe an AI invention. We do not expect that the situation could occur that, if the (initial) AI algorithm, training methods, training datasets and so forth were sufficiently disclosed, it would still be impossible for a skilled person to reproduce the invention without undue burden. Hence, a deposit system for AI algorithms does not seem to have any apparent benefits.

10.(iv)   How should data used to train an algorithm be treated for the purposes of disclosure? Should the data used to train an algorithm be disclosed or described in the patent application?

Our reply: See our answer to 10.(ii). It could be envisaged that disclosing source and format of training data, use of certain generally accessible datasets and the like may be required in order to meet the standard of sufficiency of disclosure in an AI context.

10. (v)   Should the human expertise used to select data and to train the algorithm be required to be disclosed?

Our reply: Similarly, it could be envisaged that indications may need to be provided on, for example, the expertise required to annotate certain datasets, the knowledge used in selecting data, and so forth. As an example, it may suffice to state that a skilled radiologist has annotated an image dataset by indicating whether individual images in the dataset indicate the presence or absence of a certain medical condition.

Issue 5: General policy considerations for the Patent System

11.   (i)   Should consideration be given to a sui generis system of IP rights for AI-generated inventions in order to adjust innovation incentives for AI?
*   (ii)   Is it too early to consider these questions because the impact of AI on both science and technology is still unfolding at a rapid rate and there is, at this stage, insufficient understanding of that impact or of what policy measures, if any, might be appropriate in the circumstances?*

Our reply: In our view, no new legal instruments are needed or desired. Just like in other fields of innovation, patent protection is just one option available to an innovator. An innovator may decide to go for other types of protection e.g. trade secret protection if patent protection is not considered fitting. We do not see added benefit in a sui generis system of IP rights for AI generated inventions and consider the current group of IP rights available for protecting inventions is sufficient. We are happy to be involved in future discussions on this item when impact of AI has developed further.

PHILIPS

Issue 6: Authorship and Ownership

12.   (i)   Should copyright be attributed to original literary and artistic works that are autonomously generated by AI or should a human creator be required?

(ii)   In the event copyright can be attributed to AI-generated works, in whom should the copyright vest? Should consideration be given to according a legal personality to an AI application where it creates original works autonomously, so that the copyright would vest in the personality and the personality could be governed and sold in a manner similar to a corporation?

(iii)   Should a separate sui generis system of protection (for example, one offering a reduced term of protection and other limitations, or one treating AI-generated works as performances) be envisaged for original literary and artistic works autonomously generated by AI?

Our reply: In our view, no new rules or legal instruments are required. In line with our answer to question 7, AI augments human skills and expertise but does not make human contribution any less relevant. We argued that AI algorithms cannot become inventors in their own right, as in our view they are mainly tools that support human domain knowledge and experience and are operated with human oversight. However, the output of an AI algorithm can support a patentable invention.

Following the same line of argument, if any output of an AI algorithm were considered an original literary or artistic work, such work would be considered subject to copyright.

In this case, the copyright could reside in a human being that built or designed and/or operated the AI algorithm with the aim of creating such works. The output of the AI algorithm could be regarded as a derivative work of the copyright on the original work, being the source code of the AI algorithm itself but operated with human oversight.

Thus, in line with the above, - and in line with the advice of the expert group on liability - assigning a legal personality to an AI application as a 'fictional' owner of any IP rights on the output of the algorithm is neither required nor desirable. It follows that a separate sui generis protection regime is likewise not expected to be beneficial. However, we should point out again that we would consider copyrighted works as the output of an AI algorithm to be derivative works at most.

Issue 7: infringement and exceptions

13.(i)   Should the use of the data subsisting in copyright works without authorization for machine learning constitute an infringement of copyright? If not, should an explicit exception

be made under copyright law or other relevant laws for the use of such data to train AI applications?

Our reply: As mentioned in Article 5 WIPO Copyright Treaty (WCT) and Article 10(2) WTO TRIPs Agreement, compilations of data or other material, in any form, which by reason of the selection or arrangement of their contents constitute intellectual creations, are protected as such. This protection does not extend to the data or the material itself and is without prejudice to any copyright subsisting in the data or material contained in the compilation.

Assuming that training the AI application involves copyrighted data and/or data from a copyrighted compilation, it may be argued that the use of such data for machine learning would be reserved to copyright holders or their licensees. In other words, using such data as input for a machine learning algorithm without authorization could be considered copyright infringement. In this case, the existing exception to copyright for non-commercial research or private study may apply in certain circumstances. We see no need for a separate broad exception under copyright law for the use of data to train AI applications.

13.(ii) If the use of the data subsisting in copyright works without authorization for machine learning is considered to constitute an infringement of copyright, what would be the impact on the development of AI and on the free flow of data to improve innovation in AI?

Our reply: Given the limited circumstances under which data is protected by copyright, we do not believe that this will be an impediment to the technical developments in the AI field itself. As to the free flow of data, we believe that copyright infringement by using data for AI training purposes does not impose significant additional barriers to the free flow of data that are not already in place today. This is because the existing legal framework is to be applied. Further, as already indicated under 13(i), it is believed the research exception could apply in certain circumstances wherein copyrights blocking the use of certain data for purpose of AI training would impose a significant barrier to fundamental innovation in the AI field.

13.(iii) If the use of the data subsisting in copyright works without authorization for machine learning is considered to constitute an infringement of copyright, should an exception be made for at least certain acts for limited purposes, such as the use in non-commercial user-generated works or the use for research?

Our reply: See previous answers 13.(i) and 13.(ii). We believe one could rely on the existing research exception to copyrights in jurisdictions where this provision is in place.

13. (iv) If the use of the data subsisting of copyright works without authorization for machine learning is considered to constitute an infringement of copyright, how would existing exceptions for text and data mining interact with such infringement?

Our reply: In our view, the text and data mining exception is a special form of the research exception already considered in the previous answers.

PHILIPS

13.(v)   Would any policy intervention be necessary to facilitate licensing if the unauthorized use of data subsisting in copyright works for machine learning were to be considered an infringement of copyright?

Our reply: A policy intervention is only needed if it appears that a problem exists. The existence of a problem has yet to be shown.

13.(vi)   How would the unauthorized use of data subsisting in copyright works for machine learning be detected and enforced, in particular when a large number of copyright works are created by AI?

Our reply: Detection of unauthorized use may not be straightforward for a copyright owner, though the issue is no different than for certain patents for which unauthorized use may likewise be difficult to prove. However, further to our answer to in particular question 10.(iv), if copyrighted data was used in training an AI algorithm leading to an invention for which a patent application is filed, sufficiency of disclosure requirements may require the patentee to make reference to the use of the copyrighted data in the patent application. This may be seen as an example of how copyright owners could detect unauthorized use of their copyrighted data for AI training purposes. Further, with reference to Art. 43 WTO TRIPS Agreement, it could be envisaged that under certain circumstances a party that built an AI algorithm may be required to provide evidence to substantiate that it did (or did not) use certain data for purpose of training the algorithm.

In summary, detection of such unauthorized use may not be straightforward, but it certainly does not appear to be impossible. Once unauthorized use can be shown, the existing legal framework for enforcement may be applied as is.

Again, when referring to copyright works created by AI, we point to our earlier comment that these might be considered derivative works. This could have implications if the training data used as input for the AI application itself were indeed a copyrighted work, in which case the result might be that the newly created works being output by the AI application could not be distributed without permission from the owners of the rights to the original data. At present, we cannot fully oversee whether this result would be desirable under all circumstances. This might be a topic for further investigation.

Issue 8: Deep Fakes

15.   Should the copyright system take cognizance of deep fakes and, specifically,
(i)   Since deep fakes are created on the basis of data that may be the subject of copyright, to whom should the copyright in a deep fake belong? Should there be a system of equitable remuneration for persons whose likenesses and "performances" are used in a deep fake?

Our reply: AI systems should not be allowed to create fake content. This is also in violation with principle 2 by the European Commission: *"The principle of prevention of harm. AI systems should neither cause nor exacerbate harm or otherwise adversely affect human*

beings. This entails the protection of human dignity as well as mental and physical integrity. AI systems and the environments in which they operate must be safe and secure."

In addition, whenever a deep fake is used in a criminal offense (e.g. fraud, insult), the usual provisions relating to that criminal offense should apply.

While if a movie or picture is processed to create a deep fake, copyright in that movie or picture may help, but the person affected by the deep fake is not necessarily the owner of that copyright depending on the ownership of certain moral rights. If the pictured person still holds portrait rights, it may be argued (see previous answers) that the deep fake is a derivative work and the person involved may rely on copyright enforcement to act against unauthorized distribution of the deep fake. However, generally it may be more convenient to rely on general law (e.g. tort, criminal law) rather than on IP law.

Issue 9: General Policy Issues

16.(i)   Are there seen or unforeseen consequences of copyright on bias in AI applications? Or is there a hierarchy of social policies that needs to be envisaged that would promote the preservation of the copyright system and the dignity of human creation over the encouragement of innovation in AI, or vice versa?

Our reply: Also in the field of copyright we do not see a need for new legal provisions.

Issue 10: Further Rights in Relation to Data

23.   (i)   Should IP policy consider the creation of new rights in relation to data or are current IP rights, unfair competition laws and similar protection regimes, contractual arrangements and technological measures sufficient to protect data?
*   (ii)   If new IP rights were to be considered for data, what types of data would be the subject of protection?*
*   (iii)   If new IP rights were to be considered for data, what would be the policy reasons for considering the creation of any such rights?*
*   (iv)   If new IP rights were to be considered for data, what IP rights would be appropriate, exclusive rights or rights of remuneration or both?*
*   (v)   Would any new rights be based on the inherent qualities of data (such as its commercial value) or on protection against certain forms of competition or activity in relation to certain classes of data that are deemed to be inappropriate or unfair, or on both?*
*   (vi)   How would any such rights affect the free flow of data that may be necessary for the improvement of AI, science, technology or business applications of AI?*
*   (vii)   How would any new IP rights affect or interact with other policy frameworks in relation to data, such as privacy or security?*
*   (viii)   How would any new IP rights be effectively enforced?*

Our reply: Apart from the protection offered by Article 5 WCT and Article 10(2) WTO TRIPs Agreement, users indeed can rely on sui generis database rights were applicable, and further on trade secret protection and contractual arrangements on confidentiality. It would appear that our submissions in response to the copyright questions are equally applicable to database rights. In other words, wherever a database is protected by a sui generis right instead of copyright, such right could be infringed when the protected data is used in training an AI algorithm.

In case of data that meets the requirements for trade secret protection, if misappropriation of the trade secrets occurs and the trade secret data finds its way into an AI application, the trade secret owner may rely on existing provisions to take appropriate action, for example requesting an injunction against the creator of the AI application. A further advantage of trade secret protection is that taking reasonable measures in accordance with the trade secret legal framework will generally also be beneficial in view of addressing requirements arising from data privacy and security frameworks (see 23.(vii)).

We therefore believe the existing legal frameworks are sufficient and do not see a need for any additional legal provisions.

Issue 11: Authorship and Ownership

24.   (i)   Should the law permit or require that design protection be accorded to an original design that has been produced autonomously by an AI application? If a human designer is required, should the law give indications of the way in which the human designer should be determined, or should this decision be left to private arrangements, such as corporate policy, with the possibility of judicial review by appeal in accordance with existing laws concerning disputes over authorship?
*   (ii)   Do specific legal provisions need to be introduced to govern the ownership of autonomously generated AI designs, or should ownership follow from authorship and any relevant private arrangements, such as corporate policy, concerning attribution of authorship and ownership?*

Our reply: For designs (design patents), the same holds as for invention patents: AI is a tool (like a pencil or like drawing software) operated with human oversight. The designer is the one who used the tool. No special legal provisions are needed.

Issue 12: Capacity Building

26.(i)   What policy measures in the field of IP policy might be envisaged that may contribute to the containment or the reduction in the technology gap in AI capacity? Are any such measures of a practical nature or a policy nature?

Our reply: As mentioned above, we do not see a need for new legal provisions in the AI field. There may be a new task for UNESCO, as several universities in developed countries are

PHILIPS

active in the AI field, and UNESCO may facilitate knowledge sharing between universities in developed countries and universities in developing countries.

Issue 13: Accountability for Decisions in IP Administration

*27.	(i)	Should any policy or practical measures be taken to ensure accountability for decisions made in the prosecution and administration of IP applications where those decisions are taken by AI applications (for example, the encouragement of transparency with respect to the use of AI and in relation to the technology used)?*

	*(ii)	Do any legislative changes need to be envisaged to facilitate decision-making by AI applications (for example, reviewing legislative provisions on powers and discretions of certain designated officials)?*

Our reply:	Again, we do not see a need for new rules. In administrative law, it is a general requirement for administrative decisions affecting subjects that the administrative decisions need to be well motivated so that the subject can understand the administrative decision, and if the subject appeals the administrative decision, the appeal body or court can assess the motivation. Obviously, an administrative decision cannot be motivated by merely referring to an AI algorithm as a "black box" having provided a certain outcome.