Intelligent Environments 2018
I. Chatzigiannakis et al. (Eds.)

© 2018 The authors and IOS Press.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/978-1-61499-874-7-275

Revival of the Neural Networks and the Intellectual Property Nightmare

Shubham RATHI [1]

Abstract. With the 'third wave' of Artificial Intelligence, there is a massive revival and upsurge in AI related product development. An important entity behind the AI architecture, the neural network needs to be studied carefully that adequate protection for its innovation can be secured. A key feature of a Neural Network, the Neural weights hold the inferential rules and knowledge, thus are a new way to embody knowledge and information, a new form of intellectual property to which IP laws will have to adapt. We present our discussion that sheds light on the nature of this innovation and brings to context why it is relevant to secure Intellectual Property for Neural Weights. We also rebase our arguments in the backdrop of the debates that were set off on this same topic in 1990. This paper traces the shape of this problem ever since its conception and brings to the fore the newer and expanded notions behind Neural Networks, AI and their place in the intellectual property laws.

Keywords. Artificial Intelligence, Copyright, Intellectual Property, Neural Networks, Neural Weights.

1. Introduction

The discussions around Artificial Intelligence (AI) and its tremendous impact are not new. Since almost a decade now, AI is producing output that is novel and ingenious. As this field continues to get further mainstream, as with any new technology, a lot of legal challenges are expected. For AI though, these challenges are not new. With the advent of mainstream AI in 1990s, there were massive discussions about the legal, especially the Intellectual property aspects of this radical technology by scientists, professors and legal experts. Since technology has a tendency to develop at a rate superior to the law [1], this paper takes the stride to create the technological context required for policy makers to understand and evolve the current law to fit into new dimensions that AI is evolving into. In this paper, we bring to the fore the intellectual property aspects (copyright) of Artificial Intelligence (Neural Networks) building on the discussions recorded over the last 30 years. In Section 2, we speak of the brief timeline of AI connecting it to the resurgence in this topic. In Section 3, we talk about the relevance of this discussion in the context of AI technologies that will soon beget the questions we seek to raise. Section 4, collects and builds on the arguments raised in 1990s by thinkers on the then Intellectual property aspects of Neural Networks. We base this section in the fundamental ideas of AI that are unchanged while focusing on the nitty-gritty of tech law and newer AI innovations that

[1]Corresponding Author: A 1301/2, Royale, Neelkanth Palms, G B Road, Thane (W), India; E-mail: shubham.rathi@research.iiit.ac.in

are in constant evolution. Section 5 discusses possible methods of detecting copyright infringement in Neural Networks and Section 6 concludes by throwing light on the larger inspiration of this paper, if machines are capable of 'thinking' and their legal standing.

1.1. What are Neural Networks?

The term 'Neural' is derived from the human (animal) nervous system's basic functional unit 'neuron' which are present in the brain and other parts of the human (animal) body. (Artificial) Neural Network, in general is a biologically inspired network of artificial neurons configured to perform specific tasks that traditionally can be thought as exhibiting reason. Computationally spoken, Artificial neural networks can be viewed as weighted directed graphs in which artificial neurons are nodes and directed edges with weights as connections between neuron outputs and neuron inputs [2]. The Neural Network technology is not new but has recently seen a technological uprise with the advent of Deep Learning. Neural Networks are different from computer programs by virtue of their learning style (by feeding it data), they are capable of inventive output.

1.2. What is the issue?

Since neural networks are different from conventional computer programs, there is some uncertainty about the application of intellectual property laws. One issue is the copyrightability of the set of weights: do the weights satisfy the Copyright Act's definition of a computer program and if the set of weights be said to be a work of authorship? One could argue that the network, and not a human, actually authors the weights. However, the network could also be regarded simply as a tool used by a human author, where the author chooses the data and presents it to the network. There is a confusion on how much of a Neural Network is a tool and how much is it an innovator. While humans have to be necessarily involved in the creation process but that should not make the human as its inventor. Example, If a computer scientist creates an AI to autonomously develop useful information and the AI creates a novel result in an area not foreseen by the inventor, there would be no reason for the scientist to qualify as an inventor on the AI's result. An inventor must have formed a 'definite and permanent idea of the complete and operative invention' to establish conception of the result. The scientist might have a claim to inventorship if he developed the AI to solve a particular problem, and it was foreseeable that the AI would produce a particular result [22]. Spoken precisely, if a neural network is producing output not fed to it during training and non-obvious to a person skilled in the art, which is the Neural Network is producing results outside of its domain, it ought to be considered as an innovator and not a tool. An example of this phenomenon, Alphago Zero is discussed in Section 3.2

Some aspects of Neural Network protection are well studied and caught up with law. It is widely accepted that to protect a net we need to protect three things: (i) the pattern of interconnectivity among the units, (ii) the weights on those connections, and (iii) the input and output categories, i.e., the labels that tell us what kind of numbers to put into each input [5]. The pattern of interconnectivity (the neural architecture) is rightly protected by Patents. There is little or no clarity on how the Neural Weights and the labels must be protected or whether to protect them or not since Neural weights are machine's way of embodying knowledge, a feature that the law still needs to adapt. Arguably, a ma-

chine, and not a human being, actually authors the weights in a neural network, since the human operator merely feeds data into the machine and does not know what weights (the substance behind the invention) will result after the training [6].In the context of neural networks, defining the invention is made even more difficult because of the changing nature of the invention [6] due to constant learning and updation. Since a great deal of effort may go into acquiring data and training a net, the numeric value and sequence of the weights may have considerable value, and, as a result, may be subject to unauthorized copying. The enormous investment that one might make to acquire and process data, and then to use this data to train a neural network, is all reduced to one set of easily copied weights. Accordingly, protection against theft of this valuable property is essential. This issue became quite a rage during the second wave of AI (1990s) and incited a flurry of publications and discussions around this same topic. This topic took a back foot in the later years when the second wave subsided and eventually was lost in history. Now, with the third wave of AI development, this topic is more relevant than ever before. Today we not only have Neural Networks, but its evolved version: Hierarchical Neural Networks. The role of the human has been pushed even further aback in the development stack. In 2014, Google researchers were able to demonstrate that Turing complete languages were possible using Neural Networks [17]. This research paved way for think-tanking Software 2.0, the next gen framework for writing programs composed of Neural Network weights. Microsoft is doing active work in Neural Program Synthesis where neural networks learn to synthesize programs. Naturally, the conventional copyright laws will come into question when the expression of software is not a programming language but Neural Weights. Over the years, the complexity of this issue has only become more dense. Through this paper, we hope to revive this discussion that has remained dormant for 28 years in the relevant limelight of use cases today.

2. Resurgence in the topic

The World Intellectual Property Organization (WIPO) organized a Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence in 1990 where the problematic nature of intellectual property laws for AI were discussed. In 1990 again, Lawyers wrote eloquently about the challenges posed by Neural Networks in the Intellectual Property Framework. There is very little literature relevant to this topic in the later years. The legal domain still seems to be riddled with a lot of problems spoken earlier. A peek at the historical timeline of AI connotes that this lag was because of the pace of technology. On seeing the timeline of AI and the progress of Neural Network research, it becomes evident that this downturn was because the technology had not caught up with the problems that were hypothetically posed in 1990. AI was starved of training data and what training data existed demonstrated that, depending on the architecture of the Neural Network, there would be some Neural Networks that could not be trained. Which is, the fate of Neural Networks and the problems hypothesized were not pertinent anymore.

2.1. Breakthrough Caused by Deep Learning

The application of 'Deep Learning' in neural networks was a big break through that allowed the subject to move forward. Deep Learning is a type of Machine learning that

allows a program to improve with more exposure to data and experience. As evident, the lack of quality data was a huge bottleneck in the growth of Machine Learning via the Deep Learning approach.

There have been three waves of development in the deep learning history: Deep learning known as cybernetics in the 1940s–1960s, deep learning known as connectionism in the 1980s–1990s, and the current resurgence under the name deep learning beginning in 2006 [8]. Major literature around the legal issues on the topic also emanated with these waves.

The second wave of neural networks research lasted until the mid-1990s. Funding for AI based start ups started withering when the products made were sub-par. Simultaneously, other fields of machine learning made advances [8]. In mid 1990s, deep networks were generally believed to be very difficult to train. We now know that algorithms that have existed since the 1980s work quite well, but this was not apparent circa 2006. The issue was perhaps simply that the algorithms were too computationally costly (solved by increasing model and dataset size) to allow much experimentation with the hardware available at the time [8]. This third wave of popularity of neural networks continues to the time of this writing, though the focus of deep learning research has changed dramatically within the time of this wave [8]. It is most pertinent that the legal and scientific community builds on top of the problems our predecessors of the second wave unearthed while we brainstorm and resolve the newer challenges the fast changing technology landscape poses to us.

3. Technology Landscape

As with any computer generated invention, there is often a caveat that the invention is actually computer 'assisted', to say: the role of the computer is limited to that of a tooling. This was perhaps true until a few years ago when parameters, data and even training was manual. Referring about Neural Networks, Andrej Karpathy, director of AI at Tesla goes on to state "I sometimes see people refer to neural networks as just 'another tool in your machine learning toolbox' .. Unfortunately, this interpretation completely misses the forest for the trees." [9] Neural Networks have transcended their roles as tools. They are increasingly applied in domains beyond computer science, in arts and music - domains which are classically referred to requiring creativity. Works derived out of deep creations (neural networks) are even of artistic value and so further the cause of IP protection. Neural Networks have also forayed into Software. Source code currently is protected as a literary expression under copyright. As the domain of Software Engineering evolves, we expect to see a shift in the way we write software as we transition into 'Software 2.0' where Software expression may not be literal after all. We need to encompass this change in the current Intellectual Property framework by revisiting the legal stance on copyrightability of Neural Weights.

3.1. Software 2.0

Andrej Karpathy popularized the idea of Software 2.0. Software 1.0 is the classical stack of software development as we are familiar with, written in various programming languages having basic programming operations: input, output, arithmetic, conditional, and

looping. As Andrej envisions, In contrast, Software 2.0 is written in neural network weights [9] without human intervention. Andrej mentions this as an ongoing progress in many domains: Visual Recognition, Speech Synthesis and recognition, Machine Translation, Games, Robotics and Databases. How soon or how late we are to deal with IP issues around these domains remains a speculation. Although, what seems certain is that sooner than later, the role of neural weights and the nature of IP afforded to them will have to be rethought. If the Software 1.0 written in programming languages is allowed Software copyright, do we also anticipate that its successor, written in neural weights (collection of numbers) is also capable of receiving the same perk? The answer to this seems Yes and No. The copyright law protects works "expressed in words, number, or other verbal or numerical symbols or indicia, regardless of the nature of the material objects . . . in which they are embodied Law"[18]. Which is to say that the manner of expression does not affect the copyrightability and thus protection should also extend to Neural Weights. Problems may arise when determining the Structure, Sequence and Operation (SSO) aspect of the work since Neural Weights are amorphous set of numbers that until juxtaposed on a specific neural network will mean nothing. Infact, even its developers may not be certain of its sequence or structure. For cases such as these, the existing legal regime is still to proactively think and respond to these paradoxes ever since they were was first thrown up in 1990. With the advent of Software 2.0, these questions have renewed relevance.

3.2. Alphago Zero

Intellectual Property is conferred for products that have artistic value or embody new knowledge/ creation. Software 2.0 seeks copyright for Neural Nets on the basis that they are an evolved version of the current Software regime. We can also argue that Neural weights are also in fact new knowledge. Alphago Zero is a classic example of knowledge that is harvested, learnt and applied independently by an AI system. AlphaGo Zero (AGZ), is the successor to AlphaGo, the first AI program to defeat a world champion at the ancient Chinese game of Go. Go is an ancient abstract strategy board game for two players popular in Asia. Though trivial at rules, this game is leaps and bounds more intricate than chess. Compared to chess, Go has both a larger board with more scope for play and longer games and many more alternatives to consider per move [19]. The complexity of Go is astronomical so much that number of possible moves exceeds the number of atoms in the universe [19]. AGZ may be said to be the first computer invention that in true sense fulfills the 'sweat of the brow' doctrine as unlike Alphago or even IBM's Deepmind, AGZ infers the Go rules by playing games against itself and decides on a winning strategy (self-play reinforcement learning). AGZ is not bounded by the existing knowledge/ rules of Go players. The known strategies of Go are referred to by names in language. It is hypothesized that strategies discovered by AGZ are beyond the limits of human language to express the compounded concepts [21]. This learnt language is devoid of any historical baggage that it may have accumulated over the centuries of Go study. As David Silver, of DeepMind puts it "Its more powerful than previous approaches because by not using human data, or human expertise in any fashion, we've removed the constraints of human knowledge and it is able to create knowledge itself" [20]. This 'knowledge' is in fact complex game play strategy held in the Neural Network weights. The connections that AGZ derived is knowledge, and not information simply

because prior to it, it was unknown to even the best Go players. Interesting to note here is that AGZ not only authors great strategies, but also the base training data underneath it. Simply said, Neural Networks in this particular case not only learns to connect the dots, they create the dots too and thus plead the case for copyright, again. Neural Weights for a program like AGZ are extremely valuable for the resources needed to derive them, and then for their utility. Outside the Go game, the Deepmind team is applying AGZ methods for varied problems like Protein folding [11]. If this research moves forward, soon AGZ trained Neural Weights may embody intricate knowledge about the protein biology which could be of use in cancer research.

3.3. AutoML

Google's Machine Learning platform, AutoML is a hierarchical neural network architecture that automates the process of manually designing machine learning models. In lay mans terms, this technology lets an AI build AI. As AutoML gets more mainstream and generic, the source behind the Master Neural Net's efficacy (its Neural Weights) shall be of vast value. If coupled with Software 2.0, this technology shall be among the first to demand, and then monetize its copyright over the Neural Weights.

4. Questions Raised

As discussed in the previous section, there are two forms of copyright protection that we could claim behind a neural weight. First, the software copyright that is applicable to protection and distribution of software code. Second, the intricate knowledge that the system has discovered and needs to be protected. For the latter, the most pressing question for copyrights behind Neural Weights is, if the law can recognize the intellectual creativity behind a series of apparently random numbers that not even the neural network's creator can recognize? The author opines that the law does not need to 'recognize' the creativity but rather interpret it. When the copyright office reviews a software copyright application for a source code, the jury does not dissect the code per instruction to hold up if the code genuinely does what it claims to do. The copyright application drafted in carefully worded techno-legal language helps the jury make a decision on the application. As long as the result is an original work of authorship, the copyright criteria is met. It is the inventors task here to recognize the originality and then interpret it for the application process. Similarly for the case of Neural Weights, as long as the inventor is able to establish originality of the neural weight forms, the copyright process should not be any different.

4.1. Knowledge-Information Paradigm in Neural Networks

In WIPO's 1990 symposium, Prof Thorne McCarty made a compelling example and deduction as to how knowledge was arrived at in Neural Nets. When learning the lexical disambiguation from the Brown Corpus [2], task was to construct a set of rules that will correctly classify the words in the tagged text. This task was given to a human annotator

[2] 500 naturally occurring passages tagged by hand such that every word in the text is classified in a lexical category.

and then to a Neural Network. The error rate with human was found to be 30% and the Network at 3.5%. The reason for machines superior performance was because the network internally made 12,000 rules against the 350 made by the human. Prof. McCarty notes, "From the point of view of intellectual property law, what is the valuable intellectual product here? Surely, it is 12,000 lexical disambiguation rules .. The 'knowledge', here is simply represented by a pattern of weights in the network" [14]. This statement provokes us to also ponder whether this byproduct of the neural network is knowledge or is it information? Knowledge is protected via various IPs, information on the other hand information being mere facts, is not protected. This question can perhaps be reduced to investigating if any kind of mental process or 'thinking' went behind unearthing it. The problem of speaking precisely about thought with regards to computers was identified by Alan Turing, one of the founders of computer science, who in 1950 considered the question, "Can machines think?" He found the question to be ambiguous, and the term 'think' to be unscientific in its colloquial usage. Turing decided the better question to address was whether an individual could tell the difference between responses from a computer and an individual; rather than asking whether machines 'think,' he asked whether machines could perform in the same manner as thinking entities [22]. The Neural Network certainly does not have a mind of its own to 'think' these extra rules. But its result is a product of machine thinking better known as representational learning. Prof McCarty states "Intelligent agents construct internal representations of the external world, and they process these representations in various ways to achieve their goals". Which is, for any given problem, an AI agent transforms its problem to a set of features and abstracts pattern collection, connections and thus distills knowledge which human perspective to the problem could not have achieved earlier. This byproduct hence is not mere collection of facts (information) but representational awareness or machine thinking imbibed by a network.

4.2. Dimensions of the current law

One major hurdles when copyrighting neural weights is that material from a non-human entity is not copyrightable. Section 313.2 of the U.S Copyright compendiu adds that 'The (copyright) office will not register works produced by a machine or mere mechanical process that operates randomly or automatically without any creative input or intervention from a human author'[18]. It is held here that the process is not merely mechanical (not a byproduct of only trial and error) and certainly not random. They are carefully arrived at after intensely orchestrated feature extraction and pruning.

We also need to evaluate if neural networks fall within the Copyright Act's definition of a computer program "a set of statements or instructions to be used directly or indirectly in a computer to bring about a certain result"? Does this definition adequately describe neural weights? They are certainly not "statements" in the conventional sense of the word, nor do they appear to be "instructions" both terms imply some form of sequential execution or interpretation of individual elements. Neural weights, on the other hand, cannot be taken individually; they must be taken in their entirety and, although the correct functioning of the neural network depends to a great degree on their sequence, it is not possible to predict the order in which individual weights are used [3]. Thus the SSO doctrine (Sequence, Structure, Organization) is brought into question. We can draw an analogy to the traditional software: the same way that normal software exists in two

forms, the human-readable source code and the machine-executable object code, it can be argued that the training facts are analogous to source code, while the resulting neural weights are analogous to object code. One must then contemplate the mysterious and irreversible process that connects this particular "source code" to its "object code".

Perhaps neural weights are little more than mere facts and data albeit in some arcane representational form that defies human perception. Should this then place them outside the protection of copyright law, notwithstanding their originality or the intellectual creativity needed to derive them? [3]

4.3. Neural Weights: Databases or Byte Code

We can consider Neural weights to be akin to a compilation and hence protectable as a database, or it can also be likened to Byte Code and thus considered as "object code". It could be argued, in countries like USA where databases are not protectable under copyright law, that these Neural Weights are just data. They are data only in the same way that a program written in the Java language is data, to bring Java code to life, a Java Virtual Machine 'interprets' each numeric "instruction". Thus we arrive at another conundrum: If the Neural Weights are just numbers, and Java bytecodes are just numbers, then why should Java bytecodes receive copyright protection but not Neural Weights? Both control software-implemented machine behavior. This is a logical fallacy in the law that must be addressed. There is a conceptual issue that arises repeatedly that is best expressed by the question: What is the difference between data and executable code? The answer is: it all depends on what the computer is doing with the information. For example: if a computer stores a binary file on a disk, then it's just data. If it loads that file into RAM, then it's just data. But if the contents of that file is used to control the data processing actions of the computer, then it becomes executable code (either being executed directly by the CPU or interpreted by some other software like a JVM or BASIC interpreter). Hence we arrive at the conundrum, Should weights be considered Databases (compilation of works, data or a collection of other materials arranged in a systematic or methodical way) or Byte Code (form of instruction set designed for efficient execution by a software interpreter)? Weights are most analogous to Byte Code. It is the 'instruction set' for a Neural Network but definition of Databases maps most closely to it. So it is an argument between function and form. The function of a Neural Weight is most analogous to the Byte Code but its form is most analogous to Databases. The law has to adapt to understand this conceptual shift in which we present the role of Neural Weights.

5. Enforcement Hurdles for Neural Weight Copyrights

Besides the fact that copyright for Neural Weights is far ahead of the notion that law has kept pace with and that a human inventor is necessary, there are enforcement hurdles that need to be brought to the fore. Chiefly, how does one detect and prove copyright violation? One simple answer is already well known: we could easily employ the map-maker's trick of inserting false information into the program. As road maps often carry non-existent streets, so neural nets could be trained to display the initials of the original author when given an obscure or otherwise innocuous set of inputs. Behavior like this from a competing net would give compelling evidence against independent creation [5].

However, this only prevents other people from copying the network. If another group turns out to be working independently on the same problem at the same time, the copyright obtained by the group which succeeds first has no effect on the other group, unless that other group actually copies what the first group has done.

Copyright prevents only literal copying of the network or a part of it. Although copyright is infringed by someone who copies chunks of a copyright program, it is not infringed by anyone who copies only the underlying principles or ideas to build their own version. For instance, changes of perhaps 10% seem to have little effect on the performance of the net. In view of that, protecting the exact weights is insufficient. If small variations on the initial set of weights doesn't degrade performance, how might we protect against someone who copies the original set, varies them randomly by a few percent, then claims independent creation? [5] A possible solution to this would perhaps be to embed the creators identity in the Neural Networks. A digital watermarking technology to detect intellectual property infringement of trained models was proposed in 2017 [13]. Another less intrusive plan might involve a sui generis specialized version of copyright protection for trained neural networks, perhaps one that would include the idea involved as well as its expression. Such a copyright might have a relatively short duration, say, five years. In that way, a developer could have a limited franchise for a new product without totally squelching progress [7]. It is even possible that in the future, when neural networks become so large and complex as to display reasoning powers, creativeness, and even personalities, the law will be amended to recognize them as artificial beings, in the nature of technological corporations, with separate rights and legal standing to enforce them.

6. Conclusion

This paper brings to the fore a valuable intellectual property (Neural Weights) that so far has little/ no protection. While the law takes its course in deciding the appropriate turn the policy must take, this phenomena also exposes the timeless question about the pet topic of IP scholars: how to treat output generated by an artificial intelligence. The concept of copyright dates back to the 15th century and even now most of the legal literature is derivative of the principles accepted then which did not imagine the notion of computers or their inventive output. The larger context of this paper places the question: Can computers think? (and hence create novelty) into the academic forum via the intellectual property mode. The result of this paper is intended to revive the dialogue for the need of copyrights for Neural Weights and subtly also add to the chorus of legal and scientific literature that discusses the legal status of such innovations.

Acknowledgment

The author is deeply grateful to Andy Johnson-Liard, the author of the landmark 1990 paper 'Neural Networks And The Intellectual Property Nightmare' who has been extremely gracious in discussing about this topic and fanning the author's curiosity in this domain. The author would also like to thank his supervisor, Prof. Alam, for the patient guidance, encouragement and advice he has provided throughout.

References

[1] Fenwick, M., Kaal, W.A. and Vermeulen, E.P., 2016. Regulation Tomorrow: What Happens when Technology is Faster than the Law. Am. U. Bus. L. Rev., 6, p.561.

[2] Jain, A.K., Mao, J. and Mohiuddin, K.M., 1996. Artificial neural networks: A tutorial. Computer, 29(3), pp.31-44.

[3] Johnson-Laird, A., 1990. Neural networks: The next intellectual property nightmare?. COMP. LAWYER., 7(3), pp.7-16.

[4] United States. Congress. Office of Technology Assessment, 1992. Finding a balance: computer software, intellectual property, and the challenge of technological change. Office of Technology Assessment.

[5] Davis, R., 1991. Intellectual Property and Software: The Assumptions are Broken (No. AI-M-1328). Massachusetts Inst Of Tech Cambridge Artificial Intelligence Lab.

[6] Wenskay, D.L., 1990. Intellectual property protection for neural networks. Neural networks, 3(2), pp.229-236.

[7] Graves, A., Wayne, G. and Danihelka, I., 2014. Neural turing machines. arXiv preprint arXiv:1410.5401.

[8] Goodfellow, I., Bengio, Y., Courville, A. and Bengio, Y., 2016. Deep learning (Vol. 1). Cambridge: MIT press.

[9] Karpathy, Andrej. Software 2.0 Andrej Karpathy Medium. Medium, Medium, 11 Nov. 2017, medium.com/@karpathy/software-2-0-a64152b37c35.html

[10] Perez, Carlos E. Why AlphaGo Zero Is a Quantum Leap Forward in Deep Learning. Medium, Intuition Machine, 22 Oct. 2017, medium.com/intuitionmachine/the-strange-loop-in-alphago-zeros-self-play-6e3274fcdd9f.html

[11] Kahn, Jeremy. DeepMind's Superpowerful AI Sets Its Sights on Drug Discovery. Bloomberg.com, Bloomberg, 18 Oct. 2017, www.bloomberg.com/news/articles/2017-10-18/deepmind-s-superpowerful-ai-sets-its-sights-on-drug-discovery.html

[12] Traditional Knowledge and Intellectual Property Background Brief. Traditional Knowledge and Intellectual Property Background Brief, www.wipo.int/pressroom/en/briefs/tk_ip.html

[13] Uchida, Y., Nagai, Y., Sakazawa, S. and Satoh, S.I., 2017, June. Embedding watermarks into deep neural networks. In Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval (pp. 269-277). ACM.

[14] McCarty, L.T., 1991. Artificial Intelligence And Intellectual Property Law: Some Problematical Examples. In WIPO Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence: Stanford University, Stanford (California), United States of America, March 25 to 27, 1991 (No. 698). World Intellectual Property Organization.

[15] Ronald S. Laurie, 1991. The Patentability of Artificial Intelligence Under U.S. Law. In WIPO Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence: Stanford University, Stanford (California), United States of America, March 25 to 27, 1991. World Intellectual Property Organization.

[16] Deltorn, Jean-Marc. "Deep creations: Intellectual property and the automata." Frontiers in Digital Humanities 4 (2017): 3.

[17] Graves A, Wayne G, Danihelka I. Neural turing machines. arXiv preprint arXiv:1410.5401. 2014 Oct 20.

[18] U.S. Copyright Office, Compendium Of U.S. Copyright Office Practices 101 (3d Ed. 2017)

[19] Tromp J., Farnebck G. (2007) Combinatorics of Go. In: van den Herik H.J., Ciancarini P., Donkers H.H.L.M.. (eds) Computers and Games. CG 2006. Lecture Notes in Computer Science, vol 4630. Springer, Berlin, Heidelberg

[20] Sample, I. (2017). 'It's able to create knowledge itself': Google unveils AI that learns on its own. [online] the Guardian. Available at: https://www.theguardian.com/science/2017/oct/18/its-able-to-create-knowledge-itself-google-unveils-ai-learns-all-on-its-own [Accessed 14 Apr. 2018].

[21] Perez, C. (2017). Why AlphaGo Zero is a Quantum Leap Forward in Deep Learning. [online] Medium. Available at: https://medium.com/intuitionmachine/the-strange-loop-in-alphago-zeros-self-play-6e3274fcdd9f [Accessed 14 Apr. 2018]

[22] Abbott, R., 2016. I Think, Therefore I Invent: Creative Computers and the Future of Patent Law. BCL Rev., 57, p.1079.