

Burkhard Schafer / Erica Fraser

Self-Made (Machine) Men – IP Implications of Inventions by Robots

The question if, and if so to what degree, AIs can display creativity has been an important research topic within the AI community. Innovative use of computers by artists has led to speculation of the copyright implications of computer-generated works as early as the 1950s. Less discussed however are the implications of computer creativity for patent law. This paper aims to give a first overview of the issues that the legal system may face when AIs and robots start to find novel solutions to problems which result in potentially patentable inventions.

Category: Articles

Region: Scotland

Field of law: AI & Law; Robotic; IP-Law

Citation: Burkhard Schafer / Erica Fraser, Self-Made (Machine) Men – IP Implications of Inventions by Robots, in: Jusletter IT 23 November 2017

Contents

1. Introduction
2. The Inventive Machine: Technologies and Paradigms
 - 2.1. Auto-generated Patent Applications
 - 2.2. Genetic Programming and Algorithmic Discovery
3. Patentability of Computer-Assisted Inventions
4. Beyond IP: the Self-made Machine

1. Introduction¹

«Chess is one thing, but if we get to the point computers can best humans in the arts – those splendid, millennia-old expressions of the heart and soul of human existence – then why bother existing?»

[Rz 2] What is an existential threat for Miller, the novelist, is for AI researchers a fascinating challenge. Ever since MARGARET BODEN's ground-breaking work on the conceptual foundations of computer creativity in the 1990s, the question whether AIs can be «truly» creative has been a mainstay in AI research.² Much of the academic research has tried to clarify the theoretical foundations of the field: what does it actually mean to be creative, and under what conditions would we be willing to ascribe creativity to a machine? In creative practice, artists have used robots or computers of varying degrees of autonomy for their work since the mid 1950s. Indeed a case could be made that the earliest examples of «algorithm generated art» are older still, dating back to 18th century Germany and the «composition by throwing dice» made popular by Nichols Simrock.³

[Rz 3] It was however only recently that we saw applications that not only are capable of autonomous, unsupervised creation of artistic work, but of work for which there is potentially a market. Or more precisely, we should say a market that does not care any longer if the work is computer- or human-generated. With earlier examples of computer-generated art, it was the very fact that a known (and human) artist used them to explore the limits of creativity and to challenge our conception of what it means to be an artist in the machine age that gave the work their commercial value, to the extent that they had one. The artistic merit of NICOLAS SCHÖFFER's⁴ CYSP 1 (Cybernetic Spatiodynamic Sculpture) from 1956 is that it showed how the then-resurgent «kinetic art» is an appropriate art form for a society where human existence was increasingly intertwined with machines. Similarly in 1964, Nam June Paik and Shuya Abe's Robot K-456 used robot-generated art to thematise issues of remote control and of freedom,⁵ which presupposes that the viewer knew that the work was indeed by a robot, and the intention of the human artists behind the machine. Finally, one could cite RACTER, the algorithmic poetry machine – it is unlikely that anybody would be willing to pay for its output *but for* the fact that it was created by a highly

¹ Work on this paper was supported by CREATE, the RCUK Centre for Copyright in the Creative Economy.

² BODEN, Creativity and artificial intelligence, *Art Int.* 1998, Issue 103(1), pp. 347–356; for recent discussions see e.g. VEALE, *Exploding the Creativity Myth: The Computational Foundations of Linguistic Creativity*, Bloomsbury, London, 2012 or McCORMACK/D'INVERNO (eds.), *Computers and Creativity*, Springer, Berlin 2012.

³ HAUPENTHAL, *Geschichte der Würfelmusik in Beispielen*, Dissertation, Univ. Saarbrücken, Saarbrücken 1994.

⁴ KAC, *Foundation and development of robotic art*, *Art Journal* 1994, Issue 56.3, pp. 60–67.

⁵ See BURNHAM, *Robot and Cyborg Art*, in: Burnham (ed.), *Beyond Modern Sculpture*, George Braziller New York 1968, pp. 68–77.

experimental and in this sense innovative machine, so that the proud owners of the poem holds in their hands a piece of computing history.

[Rz 4] In all these examples therefore, the artistic and/or commercial value is parasitic upon the creative vision of a human artists. We can contrast this with more recent examples of computational creativity, texts generated by programs such as Narrative Science or Automated Insights.⁶ These automated narrative generation systems take raw data, analyse it sufficiently to extract some meaning from it, and then generate natural language descriptions and explanations. An early implementation by Narrative Science was StatsMonkey, which automatically generated news stories on baseball games from raw data such as players, game score, and win probability.⁷ Subscribers to StatsMonkey are not any longer interested if the news item they receive are written by human or machine, as long as they are accurate and well-written. In contrast to the computer-generated works mentioned above, where the machine nature is constitutive for the value of the art, we call the market for these works «Turing-blind» – the more they hide their machine nature, the better.

[Rz 5] The recent emergence of «Turing-blind» markets for some computer-generated works has also rekindled the interest in copyright for computer-generated works. While speculations about their copyright implications are as old as computer art itself, the impact on practice was negligible: as long as the value of the work was crucially dependent on a known human artist using the computer to make an artistic point, assignment of authorship and exploitation rights was unproblematic, authorship resides clearly with the human controller. In Turing-blind markets, this is different. Here, we need clear answers to the question who, if anybody, owns the work that was generated by the machine.

[Rz 6] Even less studied than copyright for computer art is the question of the protection of computer-generated *inventions* through patent law -this even though this field always is «Turing-blind» in the sense described above. An invention described in a patent application either works or does not work, and is either solving a problem people are willing to pay for or not – *who* the inventor is, a famous scientists or an unknown teenager, a huge research team or a solo-amateur who struck lucky, is irrelevant for its commercial value.

[Rz 7] That nonetheless there is much less literature on robotic innovation than robot art is also due to the fact that convincing examples of autonomous machine innovation are much harder to come by. This paper tries to map out the different ways in which AIs now and in the near future will participate in the process of scientific innovation, raising some difficult legal issues in the process. We will conclude with a more speculative yet conceptually interesting question: If robots are not only going to be innovators, but innovators who can use their inventions to change themselves, will (should) this have an impact on the question of the legal status of automated devices? Are robots that «make themselves», and make themselves better, categorically different from those that remain fully determined by their human creators?

⁶ See e.g. WRIGHT, Algorithmic authors, Communications of the ACM 2015, Issue 58/11, pp. 12–14.

⁷ ALLEN ET AL., StatsMonkey: A Data-Driven Sports Narrative Writer. Proc AAAI Fall Symposium: Computational Models of Narrative, New York, ACM 2010.

2. The Inventive Machine: Technologies and Paradigms

2.1. Auto-generated Patent Applications

[Rz 8] As we saw above, one of the success stories of creative AI was the emergence of automated text generation. While traditionally seen as a problem for copyright law, these approaches become relevant for patent law when patent specifications are used as input. *Cloem* is an example of such a system.⁸ The algorithm produces a large number of permutations of a seed claim by rearranging phrases and substituting terms with alternative definitions, synonyms or antonyms. *Cloem* asserts that its algorithm is not merely random; rather, it applies patent drafting best practices to produce alternative claims. Crucially, the resulting claims are not necessarily meaning-preserving variants of the original, which would mean no «new» invention was made, just a new description of an old invention generated. Rather, they potentially enlarge the original invention's scope, or particularly in the case of substitution with antonyms, describe a distinct and new invention. These variant claims are time-stamped and optionally published online using persistent webpages. Similarly, the art project AllPriorArt.com and its sister site AllTheClaims.com use a technology that autonomously generates patent claims and descriptions. This computer system parses and randomly reassembles texts from patents and published applications from the US patent database to generate patent texts describing possible new inventions. The texts are time-stamped and published online on AllPriorArt's website.

[Rz 9] The stated purpose of both *Cloem* and AllPriorArt is to generate and publish new prior art, thereby precluding the published ideas from being patented. Publishing these texts will place the invention into the public domain before they can be patented by competitors or patent trolls. The problem with this idea is that the vast majority of claims generated by these technologies will be unintelligible nonsense. From a legal perspective, this raises the question if merely generating these derivative texts and putting them online constitutes «publication» if finding the meaningful text in between the random noise is prohibitively expensive.

[Rz 10] The patentability of an invention (or its validity in litigation) depends both on its novelty and inventiveness as measured against the state of the art.⁹ Novelty is destroyed by the existence of a prior publication that discloses what the applicant claims to have invented in a clear and unmistakable manner such that a person skilled in the art would be able to work it.¹⁰ An inventive step would be lacking where the claimed invention is obvious in light of the common general knowledge.¹¹

[Rz 11] Patent applicants are assumed to be knowledgeable of an expansive set of prior art.¹² The state of the art «comprises all matter which has at any time before the priority date of that invention been made available to the public by written or oral description, by use or in any other

⁸ See <https://www.cloem.com/flat/technology/> (all Websites accessed on 20 January 2017); *Cloem*, Untitled statement on AllPriorArt (2016) available at <https://www.facebook.com/cloempatent/posts/1689691667956972>; an example of *Cloem*-generate claims can be found at <https://www.cloem.com/media/pdf/T56231165.pdf>.

⁹ For the UK see e.g. Patents Act 1977, sections 2(1), 3.

¹⁰ *SmithKline Beecham Plc's (Paroxetine Methanesulfonate)*, Patent [2006] RPC 10; *General Tire & Rubber Company v Firestone Tyre & Rubber Company Limited*, [1972] RPC 457, at 485–486.

¹¹ Patents Act 1977, section. 3.

¹² HATTENBACH/GLUCOFT, *Patents In An Era Of Infinite Monkeys And Artificial Intelligence*, *Stan Tech L Rev* 2015, pp. 3232–3251, at 3240; *In re: Winslow*, 365 F2d 1017, 1020 (CCPA 1965).

way»¹³. Online publication on a website is recognised by the European Patent Office (EPO) and the UK courts as prior art so long as it is available without any bar of confidentiality.¹⁴ Moreover, obscurity is not a bar to a publication being considered part of the state of the art.¹⁵ The EPO Board of Appeal has held that a document stored on the Internet but only accessible via a specific URL would be considered a disclosure that had been made publically available if, before the filing or priority date of the patent or patent application, it could be found through a public web search engine by using one or more keywords related to the essence of the content of that document.

[Rz 12] Patent texts published by AllPriorArt and Cloem would likely meet this test as they *could* be found using the public search engine on their respective websites, are indexable by third party search engines, and are directly and unambiguously accessible to the public. Physical accessibility is not determinative, however, if the computer-generated texts would not actually be instructive to the skilled reader. The High Court of Justice of England and Wales has held that «available to the public» does not merely mean physically accessible, but also «sufficiently intellectually instructive [to] the skilled person using their common general knowledge.»¹⁶ In particular, a piece of art may not be instructive if «the matter may be contained in a document» but so submerged in it as not to be available.»¹⁷

[Rz 13] This is the real barrier to AllPriorArt and Cloem’s computer-generated patent texts being considered as part of the state of the art if it could be argued that the relevant prior art is hidden within vast amounts of nonsensical text. Since it would be prohibitively difficult or even mathematically impossible (depending on the size of the database) to check the auto-generated work as part of a prior art survey, it can be disregarded for this purpose. PriorArt and Cloem make important rhetorical points about deficiencies of our IP system and challenge the notion of «innovation» and «creativity», but they do not pose direct challenges to the legal regime.

[Rz 14] Are Cloem and PrioArt conceptually example of innovation by AIs? Only in the sense that the person inside Searle’s Chinese room is a speaker of Chinese.¹⁸ In that argument, Searle invites us to think of a person who is not a speaker of Chinese, inside a room. He receives through one opening cards with for him unintelligible scribbles, and then uses a rulebook to determine which card he now has to select from a stack of equally unintelligible (for him) cards, and push it through another opening. From an outside perspective, it may look as if the «room» is responding correctly to questions in Chinese. But as the speaker in the room lacks any understanding that the cards he manipulates are sentences about the external world, he obviously is not a speaker of Chinese. Even if one does not accept the general case against machine intelligence that Searle wants to make, in our case it is an apt analogy: Cloem and PrioArt only manipulate symbols, the «knowledge» they have is only knowledge about grammar and semantic relations, but no theory

¹³ Patents Act 1977, section 2(2); see also Convention on the Grant of European Patents (European Patent Convention), section 54(2): «Everything made available to the public by means of a written or oral description, by use, or in any other way, before the date of filing of the European patent application.»

¹⁴ European Patent Office, «Guidelines for Examination», sections 7.5, 7.5.1, available at <http://www.epo.org/law-practice/legal-texts/html/guidelines/e/index.htm>; *Unwired Planet International Ltd v Huawei Technologies Co Ltd & Ors* [2015] EWHC 3366 (Pat).

¹⁵ WAEDELDE ET AL, *Contemporary Intellectual Property: Law and Policy*, OUP, Oxford 2014, at 11.82; *Bristol-Myers Co’s Application*, [1969] RPC 146 T 1553/06.

¹⁶ WAEDELDE (note 15), at 11.83; *H Lundbeck A/S v Norpharma SpA*, [2011] EWHC 907 (Pat).

¹⁷ *H Lundbeck A/S v Norpharma SpA*, [2011] EWHC 907 (Pat).

¹⁸ SEARLE, *Minds, brains, and programs*, *Behavioral and Brain Sciences* 1980, Issue 3 (3), pp. 417–457.

about the external world that is utilised to come up with novel solutions to problems that are then put to a test.

2.2. Genetic Programming and Algorithmic Discovery

[Rz 15] Cloem and PriorArt take as their input human-generated text in the form of patent applications and systematically rearrange and change it. As we argued, this is not in any meaningful sense making an invention. But what if we replace human-generated text and use instead «text» generated by nature? DNA is popularly thought of as a «Code» and in this sense not different from the code that makes up our language. An algorithm that can systematically generate text using the code of natural language can equally be used to manipulate the code that is DNA, and in doing so describe new (and potentially useful) structures.¹⁹

[Rz 16] While the use of AI has assisted in creating patentable inventions for several decades, recent improvements to AI and exponential growth in computing power will likely further enable computers to produce useful inventions and become major drivers of innovation in fields like electronics, robotics, health and pharmaceuticals, materials, and nanotechnology.²⁰ Commercial services are now available that generate inventions and technology optimisations using AI algorithms in place of human intelligence. One such company, IProva, claims that hundreds of patent applications have been filed by its customers based on the inventions it has delivered, some of which have been granted.

[Rz 17] Genetic programming («GP») is a form of AI modelled after the process of biological evolution that systematically solves high-level problems by improving upon a set of candidate solutions of known performance.²¹ The algorithm creates a new generation of solutions by applying functions corresponding to genetic operations such as crossover (akin to sexual recombination), mutation and reproduction, duplication, and deletion to the best-performing solutions from the known set. This process is repeated iteratively on each generation of solutions until the software converges on a set of offspring that solves the problem, or at least sufficiently meets given termination criteria. While human operators specify the seed solutions, fitness measures, and the termination criteria, there is usually no human intervention during the program's execution²². GP has been used to independently recreate known patented inventions, generate non-infringing work-around solutions, and is responsible for the creation of at least one known patented invention that is known to have been created using GP.²³ These results have caused some to argue that GP can produce «human competitive» results²⁴.

¹⁹ SACHA/VARONA, Artificial intelligence in nanotechnology, 24 *Nanotechnology* 2013, pp. 1–13, at 1, available at: <http://iopscience.iop.org/article/10.1088/0957-4484/24/45/452002/meta>.

²⁰ NOSENGO, Can artificial intelligence create the next wonder material?, *Nature* 2016, Issue 533, pp. 22–25; ABBOT, I Think, Therefore I Invent, *BCL Rev* 2016, Issue 53, pp. 1079–1125.

²¹ See e.g. KOZA, *Genetic programming: on the programming of computers by means of natural selection* Vol 1, MIT Press, Boston 1992.

²² POLI/KOZA, Genetic Programming, in: Burke/Kendall (eds.), *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques*, Springer, Berlin 2014, pp. 145–147.

²³ KOZA, Human-Competitive Results Produced by Genetic Programming, *Genet Program Evol.* 2010, Issue 11, pp. 251–284, at 265; KOZA ET AL., Evolving Inventions, *Scientific American* 2003, pp. 52–59, at 52; KEANE/KOZA/STREETER, Apparatus for improved general-purpose PID and non-PID controllers, US Patent No US6847851 (B1).

²⁴ KOZA (note 23), at 265.

[Rz 18] In the electronics field, the antenna for the miniature satellites used in NASA's Space Technology 5 mission was designed using GP algorithms²⁵. From a set of existing antenna designs, the designers evolved a set of novel antenna designs that met the mission's predefined requirements. The unusual structural designs produced were radically different and non-intuitive, while having a number of notable functional advantages over any known antenna²⁶.

[Rz 19] So-called robot scientists are systems that integrate AI algorithms with physical laboratory robotics to autonomously conduct scientific experimentation.²⁷ Robot scientists operate with minimal human intervention to supply the consumables necessary for experiments and remove the resultant waste. Existing systems have demonstrated the ability to make observations, devise hypotheses, design experiments to test its hypotheses, employ automated laboratory equipment to run those experiments, and interpret the results.²⁸ This technology represents a marked step towards autonomous scientific discovery over the status quo where humans are primarily responsible for these functions.²⁹ One such system, «Robot Eve», has been designed for and used in drug development, particularly, identifying compounds to fight drug-resistant malaria³⁰. Given a set of 5000 molecules, Robot Eve determined the characteristics of the most effective molecules, then screened only those remaining members of the set that it predicted would be most effective. Through this process, Robot Eve «discovered» a new anti-malarial use for an existing drug that was previously known only as cancer inhibitor.

3. Patentability of Computer-Assisted Inventions

[Rz 20] What do these developments mean for the law? As AI becomes an integral tool in inventive processes, the commonly held understanding of inventiveness will be challenged. As mentioned above, patents have already been granted for inventions created using AI. To date, there has been no legislative or judicial consideration of the implications of patenting inventions generated in this manner.³¹ Given that the method of invention is not disclosed in patent prosecution, there is little reason for this to be raised as an issue as long as the process of invention does not take place very visibly in the open – we could think for instance of the above example of a new antenna «invented», as is NASA's intend, by a robot in outer space. Interestingly, US law has codified a prohibition on discriminating on this basis, declaring that «(p)atentability shall not be negated by the manner in which the invention was made.»³² Although this prohibition is not codified

²⁵ HORNBY/LOHN/INDEN, Computer-Automated Evolution of an X-Band Antenna for NASA's Space Technology 5 Mission, *Evol Comput* 2006, Issue 19, pp. 1–23, at 2.

²⁶ LOHN/HORNBY/LINDEN, An Evolved Antenna For Deployment On NASA's Space Technology 5 Mission, in: U. O'Reilly et al. (eds.), *Genetic Programming Theory and Practice II*, SSBM, Cambridge 2005, pp. 301–315, at 311.

²⁷ KING, Functional genomic hypothesis generation and experimentation by a robot scientist, *Nature* 2004, Issue 427, pp. 247–252, at 247, 251.

²⁸ KING (note 27), at 47.

²⁹ BLUNCHEN, Robot Makes Scientific Discovery All by Itself, *Wired* 2009, available at: <http://www.wired.com/2009/04/robotscientist/>.

³⁰ WILLIAMS ET AL., Cheaper faster drug development validated by the repositioning of drugs against neglected tropical diseases, *J R Soc Interface* 2015, Issue 12, p. 20141289.

³¹ HATTENBACH/GLUCOFT (note 12), p. 44; ABBOT, Hal the Inventor: Big Data and Its Use by Artificial Intelligence, in: Hamid Ekbia et al. (eds.), *Big Data Is Not a Monolith*, MIT Press, Cambridge 2016, available at: <http://ssrn.com/abstract=2565950>, at 12.

³² 35 USC §103(a).

in the UK, the results of human «search-and-test discovery» methods of invention, which are analogous with the AI algorithms described above, have to date been patentable. However, the examples that we described above do indicate that there are both doctrinal and pragmatic policy reasons why questioning the status of AI-generated inventions.

[Rz 21] The novelty of AI-generated inventions will largely depend on the inventive process used. Novelty may be absent if the algorithm used lacks variability in its outputs, or relies on similar data sets. Where an algorithm incorporates randomness or other variability, it is more likely to generate novel and distinctive inventions.³³

[Rz 22] Creativity and inventiveness require more than novelty, however.³⁴ Recall that the patentability requirement of an inventive step demands that an invention «is not obvious to the person skilled in the art, having regard to any matter which forms part of the state of the art.»³⁵ Granting monopolies over obvious inventions would contribute little to society and prevent others from engaging in technological modifications and ordinary progresses.³⁶ Given the ability of computers to supplant human intuition with brute force computational power, the notional skilled person and the bar for obviousness may need to be reinterpreted in light of AI in order for the inventive step requirement to serve its intended purpose.³⁷

[Rz 23] Assessing the existence of an inventive step concludes with determining whether «the differences between the inventive concept and the prior art constitute steps which would have been obvious to the person skilled in the art.» What will be considered obvious, and therefore non-patentable, must reflect changing inventive practices; for example, an invention that on its surface may seem inventive, may in fact be the obvious output of a computer programmed to generate inventions «like hot water from a kettle.»³⁸ This in some ways mirrors Searle's Chinese room argument: Just as we should assess the «ability to speak Chinese» from the inside perspective, not just the end result, we may have to separate for purpose of the inventive step assessment an internal perspective made in knowledge of everything that was involved in generating the idea from an outside perspective that looks at the result only. An invention that results from a computer performing a large number of trivial calculations or brute force trial-and-error testing may seem non-obvious on its face because it had not been foreseen; however, the invention may be seen as obvious because of the inevitability of discovery as anyone having ordinary skill using one of the above-described AI algorithms could have produced the same result.³⁹

[Rz 24] Here in particular, the discussion in patent law can also benefit from taking into account the experience with computer-generated art. There too a common experience is that once the mechanism by which «novelty» was created is properly understood, the impression of creativity disappears. Creativity often is in the eye of the beholder. This shifts our gaze from the origin of the work to its *identification as art*. RACTER and systems like it produce large quantities of

³³ VERTINSKY/RICE, Thinking About Thinking Machines: Implications Of Machine Inventors For Patent Law, BU J Sci & Tech L 2002, Issue 8, pp. 574–613 at 576.

³⁴ BUNDY, What is the difference between true creativity and novelty, Behav Brain Sci 1994, Issue 17, p. 533–534.

³⁵ Patents Act 1977, section 3.

³⁶ WIPO Standing Committee on the Law of Patents, «Study on Inventive Step» (2015), available at: http://www.wipo.int/edocs/mdocs/scp/en/scp_22/scp_22_presentation_inventive_step.pdf.

³⁷ VERTINSKY/RICE (note 33), p. 39; PLOTKIN, The Genie in the Machine: How Computer-Automated Inventing is Revolutionizing Law & Business, SUP, Stanford 2009, p. 102.

³⁸ VERTINSKY/RICE (note 33), p. 595.

³⁹ VERTINSKY/RICE (note 33), p. 596; PLOTKIN (note 37), p. 108.

text, the majority of it without any artistic merit. It takes human judgement to identify those outputs that have value. Consequently, just as with Duchamp's *objet trouvé*, what constitutes the artwork is not the arrangement of ink on paper (or electrons on screen) but a human singling out from the vast meaningless context one specific object and identifying it *as* art. As we saw for computer-generated patent applications, there too the majority of text is meaningless gibberish, a minority is legible text, and of these only a tiny percentage describes a new invention. Identifying these requires skill and expertise, confirming the identification through testing requires often substantial financial investment.

[Rz 25] Analysed from an economic perspective, the aim of patent law is to encourage innovation and reward the investment in research, in terms of time, effort and money. This assumes that a) innovations are scarce and b) considerable costs occur prior to an invention is being made. AI innovations as we saw challenges potentially both assumptions. They risk overloading patent offices with spurious applications and creating «aggressive» patent portfolios used by patent trolls, ultimately reducing the number of innovations that reach the market. A rational patent law under these conditions would still aim to foster research into AIs capable of creativity and innovation. But it need not automatically extend directly this protection to the outputs of such a system. Rather, it should protect the potentially significant investment, expertise and also creativity needed to identify those outputs that can be turned into a product, and in the testing, prototyping and developing that idea. A «selection theory» of creativity, in both copyright and patent law, therefore emerges as a strong contender for a legal theory of novelty worth protecting.

4. Beyond IP: the Self-made Machine

[Rz 26] We have argued above that legal protection for AI-generated innovation should be split into two aspects: first IP protection for new methods to build machines capable of innovating. Second, protection for whoever first *correctly* identifies a specific output as innovative *and* invests resources in demonstrating the correctness of their judgement. However, this solution is based on a notion of robotic invention that is essentially linear: humans design machines capable to innovate, and the machines then come up with new solutions that had not been anticipated by their designers, and therefore require (potentially) other humans to analyse this output. A potentially interesting twist happens however if the robot uses these new solutions to self-improve. In the case of the antenna design described above, we can easily imagine a robot physically changing its antenna design along similar lines, which then leads to better data that in turn drives the next wave of innovation. In this case, not just what the robot *does* can't be fully predicted by its designers, but also what it *is*.

[Rz 27] In the past, arguments that called for a radical re-think of the legal status of robotic devices and how to attribute liability to them focussed on the partial unpredictability of their behaviour. Autonomy, so the argument, breaches the causal nexus and creates as a result a responsibility gap.⁴⁰ The recent EU motion on the regulation of robotics took this idea as a starting point to push the idea of robots as holder of legal rights further than this has been done before

⁴⁰ MATTHIAS, Automaten als Träger von Rechten, Logos Verlag, Berlin 2010; БЕСК, The problem of ascribing legal responsibility in the case of robotics, AI & Society 2016, Issue 4, pp. 473–481.

outside highly speculative academic papers.⁴¹ One idea mooted in the consultation is the idea that every robot of a to-be-specified autonomy and sophistication could be associated with a unique fund, similar to the old notion of peculium from the Roman law of slavery.⁴² This money would not be controlled by owner or manufacturer of the machine and could (in part) be payment to the robot for work done. Its role would be to cover liabilities for harm inflicted by the robot to third parties similar to an insurance scheme. More radically, the proposal also suggests that it could be used as a way to compensate society for more generic harm caused by automation such as technological unemployment, similar to an income tax. The EU motion also states that they do not foresee major conceptual changes in the field of IP law.

[Rz 28] The analysis in this paper points to a different evaluation. IP, rather than delictual liability, is the field best suited to make the case for the recognition of a new sui generis status of robotic devices, both on practical and conceptual grounds. Conceptually, we can look at the foundations of property law as laid out in John Locke's labour theory of property: We own things because, and to the extent, that we mix our labour with the physical objects that we own, and thus make them an extended «part of ourselves». To the degree that robots will be able to modify and build themselves in ways unpredicted by their designers, this constitutive link for property claims is weakened. Locke's theory crucially connects self-ownership, or liberty, with the creation of property through labour.⁴³ While in his theory, the causal arrow flows from self-ownership to ownership in the things we build, what is suggested here is a reversal of this relation: mixing one's labour with oneself, improving oneself through one's effort, is constitutive for self-ownership. Pragmatically, this approach solves a problem both with the EU motion and a potential problem with our suggestion of IP protection for «identifying something as an invention». The EU motion reserves some of its more adventurous suggestions regarding the legal status of robots to only those devices that are the «most sophisticated autonomous robots».⁴⁴ It does not specify though how «sophistication» is to be measured or categorised. The capability to change itself beyond what the developer could foresee seems to be a suitable candidate for this highest degree of sophistication.

[Rz 29] Finally, as our discussion of Cloem and PrioArt has shown, robotic innovation will potentially impact severely on the administration of the patent system, with much greater numbers of patent applications, many of spurious value, being likely. The «selection model» of IP for computer-generated innovation will create additional costs to coordinate the access of multiple parties to the raw output of the robot. By attributing some of the profit from bringing computer-generated innovations to the market directly to the machine that generated them, a source of funding for these new demands could be found that allocates the costs to those players that benefit most from the development and are the main cause of the strains on the system.

[Rz 30] We thus conclude with a conditional: To the extent that a case for a legal recognition of the status of autonomous machines can be made at all (and the authors remain sceptical), robots that are capable of changing and improving themselves beyond the vision of their designers are

⁴¹ COMMITTEE ON LEGAL AFFAIRS, Draft report with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)), 31 May 2016, available at: <http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML%2BCOMPARL%2BPE-582.443%2B01%2BDOC%2BPDF%2BV0//EN>.

⁴² PAGALLO, Killers, fridges, and slaves: a legal journey in robotics, *AI & Society* 2011, Issue 26, pp. 347–354.

⁴³ For an application to IP see JUSTIN, *The Philosophy of Intellectual Property*, *Geo. LJ* 1988, Issue 77, pp. 287–366.

⁴⁴ EU Motion 2016, at p. 31 f.

the best candidates for such a change, and IP law rather than liability law the most obvious field of law to provide both conceptual-jurisprudential and pragmatic justifications.

BURKHARD SCHAFFER, Professor of Computational Legal Theory, The University of Edinburgh, SCRIPT Centre for IT and IP law Old College, Edinburgh, EH8 9YL; B.schafer@ed.ac.uk; <http://www.law.ed.ac.uk/people/burkhardschafer>.

ERICA FRASER, LLM student, The University of Edinburgh, SCRIPT Centre for IT and IP law Old College, Edinburgh, EH8 9YL; s1370756@sms.ed.ac.uk.