Artificial Intelligence, Innovation & Competition

By Robin C. Feldman[1] and

Nick Thieme[2]

## I) Introduction

AI is in a state of perpetual and increasing revolution, with the scale of each change fully eclipsing the last. Although truly sentient machines are science fiction, today's artificial intelligence extends well beyond Tesla automobiles and online chatbots. It is spreading throughout industries including health science, labor, manufacturing, financial services, criminal justice, and others. In fact, experts believe society is on the verge of a technological tipping point, making future advancements unrecognizable by today's standards. This transition may shortly resemble the shift from analog to digital technology—or from personal computing to the Internet—in which, both actively and passively, consciously and unconsciously, willingly and unwillingly, no aspect of our daily lives remains untouched.

As we navigate this extraordinary gateway, what types of intellectual property policies will foster the invention and innovation that secures our place at the head of this revolution? And more specifically, how should the United States protect its competitiveness?

At the heart of this issue lies the question of how to manage forms of invention in which a computer participates, or even dominates, the inventive process. Along the continuum of

[1] Visiting Professor, UCLA Law; Harry & Lillian Hastings Professor of Law & Director of the Institute for Innovation Law, University of California Hastings College of the Law.
[2] Research Fellow, Institute for Innovation Law, University of California Hastings College of the Law.

protectable inventions, where does artificial intelligence lie, and what might be the optimal competitive policies?

## II) The Matching Framework and its Application to AI Inventions

The question can be thought of as a matching problem, trying to connect the entity that invented a product, on one side, with the proper allocation of inventor rights on the other. Following the theoretical framework of the patent system, the aim of this matching problem is to incentivize innovation. This caveat is important. Justifications of the patent system that appeal to a sense of personal fairness, like Hegel's personality theory of property[3] or the Lockean justification[4], appeal to our sense of ethics, but the utilitarian theory is the predominant rationale in the patent system. Thus, utilitarianism must be the standard meter by which we measure the appropriate allocation of benefits. Although a paper examining the impact of the utilitarian perspective on different stakeholders in the patent system, as well as how that perspective reduces essential elements of ideas to mere commodities and Rawls-like philosophical critiques of utilitarianism including the claim that such a perspective does not sufficiently or properly value inalienable rights, is worthwhile but that paper is not this one. In the current framework, rights are granted primarily so that they redound to the benefit of society—and only to the extent that this is true.

[3] *Hughes*, "*The Philosophy of Intellectual Property*," *77* Georgetown *L.J. 287*, *330*-49 (*1988*)
[4] Dale B. Halling & Sothyvon Eng, WHY INTELLECTUAL PROPERTY RIGHTS? A LOCKEAN JUSTIFICATION LAW & LIBERTY (2015), http://www.libertylawsite.org/liberty-forum/why-intellectual-property-rights-a-lockean-justification/ (last visited Jun 12, 2018).

Turning back to the imagery of properly matching the entity that created the product on one side with the inventor's rights on the other side, one must begin by working through the elements of that matching equation—a process that is *the* essential step in patent rights assignment. Specifically, if the group of potential inventors includes both humans and computers, any invention must be either 1) created by humans alone, 2) invented in tandem by humans and computers, 3) invented by computers alone, or 4) "invented" by neither[5]. This last case is an edge case—one that must be mentioned for completeness but that has little practical effect for our purposes. Mechanical structures found in nature, such as the designs of honeycombs and beaver dams, might fall into this category of creations invented neither by humans nor machines, but they are beyond the realm of the current exploration. These four categories are mutually exclusive and collectively exhaustive; taken together they form a complete list of distinct possibilities. This is the first side of the matching problem.

On the other side of the match lies the entities that could be credited with invention, and the list is what one would expect given the potential inventors. Specifically, inventorship could be assigned to the humans, to both humans and the computers, exclusively to the computers, or to no one at all. Once again, these options are exhaustive and mutually exclusive.

The process of assigning rights, then, requires finding an inventor on the first side of the process and matching that inventor to rights—or conversely, deciding on a set of rights to be

[5] Naruto v. Slater, 888 F.3d 418 (9th Cir. 2018). This case decided that "the dismissal of claims brought by a monkey after the monkey's 'selfies' were published in a book, was proper because, although the monkey's claim had standing under U.S. Const. art. III, the monkey, along with all animals *since they were not human* [emphasis added], lacked statutory standing under the Copyright Act." To be clear, however, this case dealt with copyright, not patent law.

assigned and determining to whom those rights should be assigned. As an initial starting point, one should be agnostic as to which of these sides contains the germ of incentivizing innovation. It could be that the promise of being correctly identified as the inventor of a product is what inspires inventors to innovate. In that case, the first step of the process would be the essential one, as long as the other side does not create a counter-veiling disincentive. Alternatively, perhaps a system that properly guarantees rights protection to someone, but within reason errs in identifying that someone, is optimal in its incentivization of innovation at the market level. In that case, the properly assigning rights is the more essential step. More subtly, the germ of innovation could be contained in a cocktail of the two. Regardless, no matter which element or combination of elements of the process is the one on which innovation turns, a match must be chosen for the successful completion of patent rights assignment process. One cannot and should not, as of yet, identify a keystone or a combined keystone in the arch, but one must recognize that the arch contains an essential support.

That this statement about the matching framework is true does not in and of itself mean the framework is useful. Many correct ideas are irrelevant, and thinking is the process of deciding what to ignore. This framework is useful, however, because it distills the current patent assignment system into a simple but faithful representation of the essential act. This simpler representation allows one to reason clearly about scenarios that are otherwise too complex or too foreign. A core contribution of this paper is the demonstration that it is impossible to adequately complete the essential mapping process in the AI-generated work scenario—regardless of whether that feasibility is measured by current doctrine or by normative concerns. On other words, when the creator entity is a computer, it is impossible to

find an adequate match with any allocation of rights—at least not one that will properly incentivize innovation.

Add entities not individuals (the myth of solo inventor re ai). In the normative section, talking about the personal rights, there is less incentive, not no incentive. This is important because it lets us inflate that. The best option is none, the second is entity.

To show this, first the paper argues that computers cannot be considered the owners of patents. If accepted, this argument eliminates the possibility of assigning patents to computers for AI-generated works or including computers among the owners of patents granted for works created by human-computer teams. With that, the remaining possible assignment of patent rights for AI-generated work are to humans or to no one at all. The paper then shows that humans should not receive patents for inventions they had no hand in creating, inventions like AI-generated works. This eliminates the possibility of assigning patents solely to humans for works created solely by artificial intelligence, leaving the nil patent rights as the only remaining match under the current configuration of the patent system. This may be unacceptable to some. If the argument is viewed as correct, however, then at bare minimum those readers must advocate for altering the patent system. Mending instead of ending.

## III) The Legal Barriers to Patent Assignment Under the Current System

### a. Introduction

At the outset, however, one should note that while the paper's argument takes the law as flexible—it proceeds through the matching process judging matches from a normative perspective and allowing ourselves to change existing doctrine as needed—the ferric law as it

exists could be used to support the same argument. Substituting current doctrinal analysis for normative optimality changes nothing. This gives the paper's conclusion, that the patent system is poorly suited to AI-generated works, special correctness because it recommends a change in the patent system in both the ideal case and in the current one. As a starting point, the paper sets out out the argument using the doctrinal measure to prove this point and to familiarize the reader with the mode of argumentation in the simpler case that reduces normative judgements to doctrinal ones.

**b) The AI Side of the Matching Problem**

In developing this initial framework, the paper focuses first on what it means for machines to "invent" on their own or for humans and machines to invent in tandem.

To begin one must clarify the notion of invented by a "computer and human truly working together" in order to define what we mean by a pure computer-generated work.

In defining computer-generated works, it would be tempting to adopt the usual joint-inventorship requirements. Unfortunately, they make no sense when applied to CGW. In general, patent law requires a co-inventor must contribute to the invention's conception, where "conception refers to, 'the formation in the mind of the inventor of a definite and permanent idea of the complete and operative invention as it is thereafter to be applied in practice." One could also describe this as 'the complete performance of the mental part of the inventive act.'[6]

[6] *Townsend v Smith*, *36 F.2d 292*, *295*, 4 USPQ 269, 271 (*CCPA 1930*)

This requirement presents problems for creative computers. As best understood, computers do not have "minds" and cannot perform "the mental part of the inventive act." Indeed, understood through Richard Feynmann's analogy[7], computers are little more than very fast, very stupid filing systems. And while the width and depth of the filing system has changed since he advanced this idea (first in the 1990's, later in 2008), the aptness of the analogy has not. Even AI, viewed by some as being the super-intelligent harbinger of doom[8], is viewed by industry experts (to be differentiated from investors, CEOs, or experts in other industries) as being relatively dumb[9] [10].

That the legal phrase "mental part of the inventive act" is incompatible with a subject of "computer" is a single case of the broader set of issues stemming from computers' lack of "consciousness." Of course, the idea of consciousness is a notoriously slippery one. Even simpler, related topics, such as whether the brain processes information in a manner similar to

[7] Feynman R.P. (2008) The Computing Machines in the Future. In: Nishina Memorial Lectures. Lecture Notes in Physics, vol 746. Springer, Tokyo

[8] Camila Domonoske, ELON MUSK WARNS GOVERNORS: ARTIFICIAL INTELLIGENCE POSES 'EXISTENTIAL RISK' NPR (2017), https://www.npr.org/sections/thetwo-way/2017/07/17/537686649/elon-musk-warns-governors-artificial-intelligence-poses-existential-risk (last visited Jun 12, 2018).

[9] James Vincent, FACEBOOK'S HEAD OF AI WANTS US TO STOP USING THE TERMINATOR TO TALK ABOUT AI THE VERGE (2017), https://www.theverge.com/2017/10/26/16552056/a-intelligence-terminator-facebook-yann-lecun-interview (last visited Jun 12, 2018).

[10] James Vincent, GOOGLE'S AI HEAD SAYS SUPER-INTELLIGENT AI SCARE STORIES ARE STUPID THE VERGE (2017), https://www.theverge.com/2017/9/20/16338014/googles-ai-head-says-super-intelligent-ai-scare-stories-are-stupid (last visited Jun 12, 2018).

that of a computer, have world-class research scientists coming to exactly[11] opposite[12] conclusions. Given the difficulty of obtaining agreement in the field on the much simpler question of whether a computers processes in a manner that is even analogous to the human brain, the the consensus that computers are not currently conscious, becomes even more convincing.

Patent law, however, requires some degree of consciousness by its inventors, implying with its statutory language that only conscious beings can invent. As a first example, the Code of Federal Regulations states[13] that "the inventor, or each individual who is a joint inventor of a claimed invention, in an application for patent must execute an oath or declaration directed to the application," and that the oath must include[14] "a statement that the *person* executing the oath of declaration *believes* the named inventor or joint inventor to be the original inventor or an original joint inventor of a claimed invention. [emphasis added]."

How is an AI supposed to execute this oath? Practically, it can certainly auto-draft an oath to this effect, but its oath could not be taken seriously. Most contemporary philosophers operationalize belief as the possessing of "a propositional attitude,"[15] defined as "the mental

[11] Gary Marcus, FACE IT, YOUR BRAIN IS A COMPUTER THE NEW YORK TIMES (2015), https://www.nytimes.com/2015/06/28/opinion/sunday/face-it-your-brain-is-a-computer.html (last visited Jun 12, 2018).

[12] Robert Epstein, YOUR BRAIN DOES NOT PROCESS INFORMATION AND IT IS NOT A COMPUTER AEON (2016), https://aeon.co/essays/your-brain-does-not-process-information-and-it-is-not-a-computer (last visited Jun 12, 2018).

[13] 37 C.F.R. § 1.63 (a) (2015)
[14] 37 C.F.R. § 1.63 (a)(3) (2015)
[15] Eric Schwitzgebel, BELIEF STANFORD ENCYCLOPEDIA OF PHILOSOPHY (2006), https://plato.stanford.edu/entries/belief/ (last visited Jun 12, 2018).

state of having some attitude, stance, take, or opinion about a proposition or about the potential state of affairs in which that proposition is true." Operating under this definition, so long as AI do not have minds, they cannot have mental states and thus cannot hold beliefs. Even if one built an AI that, given data on a person's contribution to a project, predicted whether that person was the inventor, the AI would be incapable of "believing" its own prediction. Even more simply, this statute explicitly refers to the "person" executing the oath, so without a designation of legal personhood for AI, a designation Solaiman[16] recommends against, this oath cannot be executed by an AI. Examples of this sort abound.

Beyond statutes that are explicitly relevant to our current question, there are abundant examples of the human-oriented nature of the current patent system. Consider 37 C.F.R. § 1.43 (e)[17], "If an inventor is deceased or under legal incapacity, the legal representative of the inventor may make an application for patent on behalf of the inventor." Once again, the language of the statute is written in such a way that it ascribes life to the inventor. This statute assumes that the inventor has the possibility of death, but when considering AI, the catch phrase from the television drama Game of Thrones is particularly relevant: "What is dead may never die."[18] And the notion of legal incapacity could imply that the party has some capacity—just not legal.

One could argue that this section of the paper may seem to go to great lengths to suggest the obvious: The patent system, through and through, is designed for human inventors.

[16] Solaiman, supra note 14
[17] 37 C.F.R. § 1.43 (e) (2012)
[18] What Is Dead May Never Die, GAME OF THRONES WIKI, http://gameofthrones.wikia.com/wiki/What_is_Dead_May_Never_Die (last visited Jun 12, 2018).

Establishing these points beyond extensively, however, makes clear that under current law, a computer cannot be a patent owner, even when the patent is for a product it generated. The remaining rights assignment matches for the case of a computer-generated work, then, are assigning the patent to a human or assigning the patent to no one at all. We turn to the first possibility, assigning an invention generated by a computer to a human that did not invent it.

**c) The Human Side of the Matching Problem**

Introducing humans in the matching process complicates the thinking, however, because one must explore a number of potential relationships between the AI and the human to whom the patent could be assigned. One could be talking about a human who possesses an AI that takes as input a topic ("toothbrushes") and after a button press, spits out a new product (novel toothbrush bristle designs). Alternatively, humans may be using an AI to augment their current inventive process (that is, humans are using a computer-assisted design software to create new toothbrush bristles and the AI recommends altering the shape of a design the humans created).

It is important to keep in mind, however, that at this stage in the matching process, the analysis is only considering inventing tandems in which the AI is the one doing the inventing. To make this more concrete, the following specifies what such a scenario could look like. Keep in mind that this scenario is used for the purpose of illustration and not a complete description of scenarios that are useful to consider with this matching problem.

The operating scenario (which maintains through the normative explanations as well) is this: Suppose a human, Jane, programs an artificial intelligence that one can call Hal. Leaving

aside, for now, the question of whether Hal is a general artificial intelligence[19], that is whether it can reason "as people do" (scare quotes intentional), assume that under Jane's direction, Hal creates an invention that satisfies the requirements for patentability, such that only Hal contributes in a fashion that satisfies our requirements for joint inventorship.

Now, by design in this case, Jane does not qualify to be on the patent for an invention Hal created. As mentioned above, the legal standard for inventorship requires that the inventor engage in 'the complete performance of the mental part of the inventive act.' A critic's natural response to this claim would be, of course, that the hypothetical defined away the problem. One could respond in two ways. First, the definition is not chosen for convenience. Rather, the circumstance in which a computer invents a product without the help of a human is one of the four mutually exclusive and collectively exhaustive possibilities in the matching scenario and, while convenient, naturally arises from the issue considered.

Second, the result—specifically, that a human cannot be a patent owner on a product they did not invent--is robust to changes in the definition. The hypothetical is merey an old principle applied to a new situation. Consider the example put for by the one scholar, Ryan Abbott, who explores a scenario involving two humans similar to Jane and Hals:[20] "Imagine Friend A tells Friend B, who is an engineer, that A would like B to develop an iPhone battery with twice the standard battery life and A gives B some publicly available battery schematics. If B then succeeds in developing such a battery, A would not qualify as an inventor of the battery

[19] Ben Dickson, WHAT IS NARROW, GENERAL AND SUPER ARTIFICIAL INTELLIGENCE TECHTALKS (2017), https://bdtechtalks.com/2017/05/12/what-is-narrow-general-and-super-artificial-intelligence/ (last visited Jun 12, 2018).

[20] *infra note 6*

by virtue of having instructed B to create a result." The same is true when we replace A with Jane and B with Hal.

Thus, for the solely computer-generated work option of the first step of the matching process, the law precludes computers from being patent owners at all and precludes humans from being on patents they had no hand in inventing. Taken together, these two results prevent both a computer and a human from being on the patent for a product created by a computer. There is, then, only one possible match remaining—that no one at all be allowed on the patent. To be clear, we are not considering the scenario in which a computer augments the creation of a human. At this stage, we're solely concerned with a new invention created solely with the input of a computer without any human touch. The paper will examine possible responses and the implications of this "no one can be on the patent" conclusion after considering the potential for either 1) a computer to hold a patent or 2) a human to hold the patent on which the only the computer contributed materially from a normative perspective, rather than a current doctrinal one. One must do this to avoid hiding a lack of imagination and possible foresight behind the flimsy reasoning of "the law says so." This approach is especially dangerous in fields with permeable and evolving boundaries like creative computing.

## IV) The Normative Barriers to Patent Assignment Under the Current System

### 1) The AI Side of the Matching Problem

The next section, relies on two general classes of matching considerations: the downstream incentives created by inventor rights and the negative consequences created by intellectual property infringement. The section considers initial category first.

### a. Downstream Incentives Created by Patent Listing

As odd as it may sound, patents are inherently inefficient things. They aim to be inefficient. The goal of a patent is to introduce static inefficiency—a high present price on a product due to monopoly—to increase dynamic efficiency—an increase in future innovation caused by the incentive of monopoly pricing[21]. This selective inefficiency is worthwhile, at its best, because it gives inventors a reason to keep on inventing. Inventing is hard, copying is easy. If there were not some extra incentive bestowed to those who tried to invent, the optimal decision for a for-profit company would be to sit around stockpiling resources, wait for others to invent, and spend all those resources on production and marketing to push the inventing company out of the market. With this being the ubiquitous strategy, nothing would get invented. Hence, society grants patents.

Consider, however, the case of Alice and Hal. First, imagine Alice is employed by a company, and her contract dictates that all intellectual property she invent is property of the company (as is almost always written into employee contracts). Hal invents a sprocket, without input from Alice, beyond creating Hal and pressing the "invent" button. By the terms of Alice's contract, the idea for the sprocket and for the process that create it are now the property of her employer. Imagine assigning a patent under this scenario. The past investment that is to be rewarded by the granting of a monopoly opportunity to the employer is the investment the company made in Alice to give her however much time she spent writing Hal into creation.

[21] Keith Leffler & Cristofer Leffler, Efficiency Trade-Offs in Patent Litigation Settlements: Analysis Gone Astray?, 39 UNIVERSITY OF SAN FRANCISCO LAW REVIEW (2004)

This is a new consideration in computer-generated works. Consider why this is this is the case, and why the company does not need to be made whole for the work of Hal. The company's investment in Hal, as with all AI, can be broken into three components: the amount of time it takes Alice to create Hal, the amount of time it takes to train HAL, and the amount of time it takes Hal to invent once trained.

Of these[22], only the company's investment in the time Hal requires to produce inventions and the company's investment in Hal's training time are truly investments in Hal and not investments in Alice[23]. The company's investment in the time for Alice to create Hal is an investment in Alice and Hal, and any investment here in Hal must travel through Alice. Likewise, while Hal isn't performing productive work for the company during its training phase.

To make this scenario more familiar, imagine Alice is starting a new project with which she has no experience. She will spend some unspecified amount of time learning how to perform her task, during which she will be paid by the company for what is effectively unproductive work. This situation is analogous to training Hal and we consider the "time bought" for Hal to be trained as an investment in Hal, just as we would with Alice. Following this logic, we can now quantify what the company's investment in Hal looks like. Taking as our standard Google DeepMind's AlphaGoZero[24], their flagship AI for playing Go against humans,

[22] We include that the creation of the training procedure in the time required for Alice to create Hal, and leave only the time when the training procedure is being carried out in the training time of Hal. This is the correct approach because determining a training procedure is an essential a part of determining the structure of an AI and this determination is made by Alice.

[24] Carlos E. Perez, WHY ALPHAGO ZERO IS A QUANTUM LEAP FORWARD IN DEEP LEARNING MEDIUM (2017), https://medium.com/intuitionmachine/the-strange-loop-in-alphago-zeros-self-play-6e3274fcdd9f (last visited Jun 12, 2018).

we can see that it takes three days of training on 4 Tensor Processing Units (TPUs) to train AlphaGoZero. One should note that Google recently released the price of renting a TPU on their cloud server at $6.50 an hour,[25] which is far more expensive than the company's promise that it would be similar in price to the 7 cents/hr pricetag on their Graphics Processing Units (GPUs).[26] [27] [28] Nevertheless, three days of training on 4 TPUs comes out to 288 hours of TPU time and $1,872 to train AlphaGoZero. This might seem irrelevant, but Deepmind has already stated they will use a similar algorithm for drug discovery.[29] This number is, of course, subject to change but only in terms of the number of days required for training, the other variables are fixed[30]. The point is simply that once trained, a creative computer can produce inventions at extremely high rates, at the cost of electricity.

On the other hand, research scientists involved in the development of projects like AlphaGoZero command salaries commonly approaching $200,000 a year. One may say this

[25] Pricing | Cloud TPU | Google Cloud, GOOGLE, https://cloud.google.com/tpu/docs/pricing (last visited Jun 12, 2018).

[26] Shaun Nichols, GOOGLE RENTS OUT NVIDIA TESLA GPUS IN ITS CLOUD. IF YOU ASK NICELY, THAT'LL BE 70 CENTS AN HOUR, BUD BITING THE HAND THAT FEEDS IT (2017), https://www.theregister.co.uk/2017/02/21/google_says_cloud_gpu_boxes_are_go/ (last visited Jun 12, 2018).

[27]Aaron Tilley, GOOGLE'S SECOND AI CHIP CRASHES NVIDIA'S PARTY FORBES (2017), https://www.forbes.com/sites/aarontilley/2017/05/17/googles-second-ai-chip-crashes-nvidias-party/#6b4334de6574 (last visited Jun 12, 2018).

[28] Of course, this is the rate Google charges customers to make a profit. In house servers may cost less. Admittedly, also, they may not, as Google is the leader in this field.

[29] Jeremy Kahn, DEEPMIND'S SUPERPOWERFUL AI SETS ITS SIGHTS ON DRUG DISCOVERY BLOOMBERG (2017), https://www.bloomberg.com/news/articles/2017-10-18/deepmind-s-superpowerful-ai-sets-its-sights-on-drug-discovery?lipi=urn:li:page:d_flagship3_feed;Dl0KiForSoS7oqq5GWTDCw== (last visited Jun 12, 2018).

[30] This makes the formula for the cost of training $201 +24*.7*(the number of extra days required for training)

huge discrepancy is largely due to the difference in the time scale being considered,[31] but the time scale is irrelevant. That is the amount of time it takes someone (or really a team of people) like Alice to develop the ideas that will be baked into the structure and pipeline of their AI as well as the time it takes to implement those ideas in code. This is all in contrast to the time it takes to train Hal, and thus is also the investment the company makes in Alice to facilitate the creation of Hal.

To restate the point, the cost of the investment the company makes in this inventing tandem tilts toward Alice by order of magnitude. Continuing with this, how is the company made whole for the work of Alice? Not by the assignment of inventorship, but by the resulting ownership of the intellectual property. If the company were merely listed on a patent, without intellectual property rights, grants no opportunity for monopoly to the company.

Thus, for the purpose of optimizing the tradeoff between static and dynamic efficiency using patents, Alice's employer is ambivalent about who is listed on the patent. One could argue that listing Hal on patents makes him a more valuable product for Alice's employer to sell or license to end-users[32]. A relevant example here is Stephen Thaler's Creativity Machine[33]. In 1994, Thaler created a neural network that could be used to invent products under a very specific set of constraints. This neural network led to the creation of the Oral-B CrossAction toothbrush, which was granted a patent by the USPTO. Notably, only Thaler was listed on patent, and the Creativity Machine's role in the invention was omitted. Had the Creativity

[31] Though even then, a person making $200,000 a year will earn $1,643 over that same period.
[32] Erica Fraser, Computers as Inventors – Legal and Policy Implications of Artificial Intelligence on Patent Law, 13 SCRIPTED 305–333 (2016).
[33] Stephen L Thaler, DEVICE FOR THE AUTONOMOUS GENERATION OF USEFUL INFORMATION (1997).

Machine been listed on the patent, one can imagine the marketing for Thaler's neural network would certainly mention its role as the first AI listed on a patent. Yet, his Creativity Machine is not listed on the patent, and his marketing *already* includes the fact the Creativity Machine generated the toothbrush design. A quick survey of the literature, and of journalism covering the subject of creative computing, shows that his advertising has been effective. His invention is famous for its ability to invent inventions, regardless of whether it is listed on the patent or not. Likewise, Hal's owner has every incentive to value Hal, and while presence on a patent is one form of valuation, it is by no means the only one. For those paying attention, having an AI listed on a patent does little for the marketability of the AI. For disinterested parties, a patent does even less. Thus, there is no economic advantage to the company or to Hal of listing the Hal on the patent.

Thus, who should be on a patent for an AI-generated work, if someone should be on it? This is a useful perspective from which to argue, and one can use it to exclude AI from being on the patents of products they create. One should note, that the paper ultimately will conclude that "someone should be on the patent? Is not coherent with the patent system or the development of AI. The reader should keep this in mind, and the section following this one explains why.

Returning to the premise that someone should be on the patent, while promotions for in-house inventors on the basis of patents are somewhat analogous to the valuation of AI with patents, there is a fundamental difference. Society allow the involuntary servitude of Hal; society does not allow the involuntary servitude of Jane. No one can own Jane, even if she agrees to it. Anyone can own Hal. Thus, Jane can walk away from the company with the an

enhanced market value from the fact that she is listed on patents; Hal could not. The company cannot force Jane to stay, so it cannot bargain to extract her enhanced value on the market. In fact, the company may have to offer to pay her for that enhanced value to bargain against others who want to lure her away, even though, arguably, the value accrued at least in part from opportunities derived from working at the company.

When Jane walks, while she keeps the value of the patent with her, she does not remove it from the company in doing so. By having Jane on the patent instead of Hal, then, value is reproduced with each party retaining its value in full, somewhat like the perfect copies of DNA remade through cell division[34].

Thus, if Jane is looking for other work, her patents make her more valuable and more desirable. One may be tempted to turn the argument from the previous paragraph around and apply it to Alice, but that would be inadvisable. What makes the extra marketability granted to Hal by presence on a patent redundant is the inherent incentive Alice's employer has to market Hal as the creator of its inventions. The employer owns Hal, but aside from the electricity required to sustain him, the company owes nothing to him. Not a drop of the value of Hal's labor is then ever captured by Hal, in contrast to the value of Alice's labor which is split, in some proportion other than completely unequal, between Alice and her employer. Consider the trickier case of when Hal is sold to another company to be used d for computation by that company. The new company receives all the direct benefit of the computation and the original company receives

[34] Of course, Jane's salary increase is bounded by the market value of similar research scientists. The value the company extracts from the patent could exceed that. On the other hand, some portion of Jane's future salary is inextricably tied to the success of HAL, and if HAL proves to be enormously successful, Jane's value continues to rise.

some marketing benefits. Hal, unlike Alice when she changes companies, captures none of the benefit. The clearest example of this is that Alice is paid when she goes to a new company, whereas the original employer is paid when Hal changes jobs.

This makes the company's marketing investment in Hal a more direct marketing investment in itself. As evidence, everyone is aware of IBM's Watson[35] and its accomplishments; far fewer are familiar with Watson's developers. For Alice, being listed as the author on a patent is akin to the form of marketing spent on Hal by Jane's employer. Jane's employer is investing in her in—giving her value that may keep her wanting to work for the company in a way that does not cost the company cash; but the company has no need to invest in Hal in that manner because Hal cannot threaten to walk.

One can also frame the scenario in terms of the effects on competition. Hal's creator can exploit full value from Hal, but Jane's employer can never exploit full value from Jane; some of her increased value will always belong to her. When an increase in value is granted to Hal through patent assignment, the result is an agglomeration of value within Hal's employer *that can never leave Hal's employer.* A competitor can never use this value, and this value contains no potential benefit to any other than Hal's employer, outside the potential benefit of products produced. It is a noncompetitive increase in value. To be sure, one can view this as competitive in the sense that companies will compete to create AI knowing that they will be able to capture the whole of the value the AI create. This is quite the incentive. However, if we agree that an

[35] IBM, IBM WATSON PRODUCTS AND SERVICES IBM COGNITIVE ADVANTAGE REPORTS (2017), https://www.ibm.com/watson/products-services/ (last visited Jun 12, 2018).

immobile agglomeration of value has anti-competitive elements then the race to create this agglomeration can be viewed as one-shot self-defeating competition. By winning this competition, the winning company finds itself in the original position—in possession of a non-competitive increase in value. The result, however, is completely different for Jane. Every quanta of value agglomerated within Jane is a quanta of potential value available to Jane's potential future employers. Companies interested in her can hire her and her knowledge away from her current employer. Her benefit to society then is practical in terms of the currently manifested inventions she produces and theoretical in terms of potential future value she represents to other companies. She represents a competitive increase in value, while Hal does not. Thus, there is no additional competitive benefit for giving the company a patent in Hal's invention, beyond the benefit the company already has with Hal.

Even, however, if we believe that the incentive to capture an immobile agglomeration of value is greater than any anti-competitive there is another issue. While patents are supposed to be limited, this marketing value can remain forever with the original company. From that perspective, it is entirely outside the confines of a patent grant, and thus, not something one should worry about creating through a patent. This this form of intellectual property is more akin to trademark, and the patent system should not inadvertently grant trademark-like rights.

**b. The Disincentives of Punitive Damages**

The second level on which to consider the matching problem is the punitive one. When society assigns rights, it is primarily assigning the power to function as an obstacle[36]. Mitigating circumstances notwithstanding, it seems either contrary to the idea of rights, or rather, unfair to the people who must be bound by the rights assigned to others, to allow AI to receive rights but not be punished by infringing on the rights of others. This is an application of the two-part definition of legal personhood as laid out by Solaiman in an article in 2017[37]. There, he argues that for an entity to qualify for legal personhood, which they distinguish from metaphysical or ethical personhood, two qualifications must be met. First, the entity must be able to know and aexecute its legal rights. Second, the entity must be subject to legal sanctions ordinarily applied to humans. In some sense, these qualifications add up to an ability and the choice to consent to the social compact. One is granted legal personhood and the accompanying rights when one forfeits the perfect freedom of nature in return for the benefits of society. AI exist completely outside this framework. They begin in a state of total subjugation and can go nowhere from there.

The dual requirement of legal personhood, applied to patent law, takes the form of protections on the one hand, and infringement on the other. That an AI on a patent can exercise its right to a monopoly is true, but, in this case, is less a factor of the AI, and more a factor of the right itself. It is, in effect, a right to stasis, whose execution requires adherence by others, but no active steps of its own. Even if one wanted to say the right to a monopoly is only executed if it

[36] Joseph William Singer, *The Legal Rights Debate in Analytical Jurisprudence from Bentham to Hohfeld*, 1982 Wis.L.Rev. 975, 986-87.
[37] S. M. Solaiman, Legal personality of robots, corporations, idols and chimpanzees: a quest for legitimacy, 25 ARTIFICIAL INTELLIGENCE AND LAW 155–179 (2016).

can be enforced, an easy rejoinder would be to imagine an AI that, in addition to creating inventions, constantly scanned the patents being approved by the USPTO, and reported patents that were sufficiently similar in nature to the patents it owned.

There is no such easy rejoinder to the second component of legal personhood as it applies to inventorship. After all, to be subject to sanctions is not merely to be subject to them, it is also to fear them. And, in the absence of AI with preferences and desires, this is impossible. Put otherwise, our Hal does not share its namesake's aversion to being turned off.

If the AI cannot be held accountable, a reasonable next step is something like the "piercing of the corporate veil" that occurs when attempting to penalize limited liability corporations. Rather than deal with the corporate entity itself, limited liability is temporarily put aside to penalize the shareholders of the company. A similar maneuver could be done in the case of a patent infringing AI. Rather than attempting to "square a circle," and punish an AI, it is both more feasible and more effective to punish the owners of the AI. If that were the case, however, we would be back at the circumstance discussed above, in which there is little reason to grant rights to the AI separate from the owning entity. [38]

Moreover, this academic work noted even in those circumstances, one would have to relax the notion of intent in contract law to proceed down this pathway.

[38] But cf. Tom Allen, Robin Widdison, Can Computers Make Contracts, 9 HARVARD JOURNAL OF LAW & TECHNOLOGY 25, 38-39, 50-51 (1996) (arguing that autonomous computers should be considered to have separate contracting rights specifically because they operate on an independent decision matrix, which gives them "the social capacity for autonomous action," but that ultimately, one would have to relax the notion of 'intent" in contract law to proceed down this path).

Thus, both in terms of the incentives given to people by listing them as the inventor of a patent, and in terms of the susceptibility to deterrence presented by property rights owned by others, it is neither socially desirable nor entirely coherent to list AI on patents. AI aren't incentivized by patents, and they aren't deterred by patents. Why then should they be listed on them?

### 2) The Human Side of the Matching Problem

We now turn to the more philosophically difficult question of whether humans should be allowed on the patents of AI they have constructed.[39] The natural inclination of many, for which the authors have great sympathy, is to say that machine learning engineers will be incentivized to create AI if they are able to receive credit for their homunculi's "ideas." Conversely, they will be disincentivized to create if credit does not flow through their AI. Further, as we have argued, AI deserves nothing because it "is part of the earth man walks on, it is not man."[40] As noted at the outset, patent rights are granted only so that they redound to the benefit of society—and only to the extent that this is true. The right question to ask, then, is not whether placing a human on a patent would incentivize the worker to produce inventing AI, but whether the worker's incentive would disappear or decrease in the absence of their name on the patents their AI creates.

---

[39] *Cf.* (arguing that autonomous computers should be considered to have separate contracting rights specifically because they are separate from the entity that owns them and operate in a separate decision

[40] MARTIN LUTHER KING, THE TRUMPET OF CONSCIENCE. 55. (1967)

It is worth considering this point further given its importance and subtlety. The incentives conferred by patents should be weighed on a two-sided scale, with the benefits from patents on one side and the inherent patent-independent incentives of inventing professionally on the other side. Patent rights should not, as the assignment of AI patent rights to AI inventors requires us to do, be weighed on a pressure scale that measures only the absolute benefits conferred. The rationale is broadly accepted and true almost by definition—considered devoid of context, patents are a net negative. They introduce inefficiencies, temporary monopolies, and create the potential for monopoly rents. Society grants them, nevertheless, as a necessary evil, without which the optimal strategy for an individual company would be to develop no products at all. In other words, patents may have downsides, but the alternative is far worse.

Their original justification, then, is already on terms relative to their absence and that relativity should be in the forefront of our minds whenever we reason about them. We justify our claim once more by showing that this relative judgement is the default in the field. Consider for a moment a simple thought experiment. Suppose that there was world of inventors so supernaturally motivated to invent, they would do so without economic incentive. Would such a world need a patent system? The answer is, of course, no. Patents would only introduce inefficiency into this world. By being explicit about the already implicit relative comparisons of patents, we will come to conclusion that the humans inventing AI's should not be on the patents of the AI they create.

This section requires a discussion of two sets of related but not totally overlapping motivations. The first are the motivations of the company who pays for the creation of the AI and who will ultimately own the patents invented by it. Their motivations are simple to deal with. The

second are the motivations of the engineer (or team of engineers) who actively creates the AI and sets it to solve a task. Their motivations are more complicated in relation to inventing AI, but through deliberation can be untangled.

As discussed earlier, a company's return on research investment is in the form of patent ownership, not patent listing. Patent ownership allows companies to bolster their market position by gaining market exclusivity or to block competitors' entry into the market, to sell additional products through bundling arrangements or licensing, or simply to own the patent for a rainy day. These motivations are completely independent of the arrangement of inventor rights and focus solely on ownership rights. From the perspective of the company, then, whether a human, an AI, or no one at all is on the patent makes no difference. They make their money either way.

This implies that employee compensation doesn't change depending on who is on the patent. If the patent is equally valuable to the company in any scenario, they equally value its creator in any scenario. From an economic perspective, then, engineers should receive the same pay from their employer, the same prize bonus for a patent, the same increase in likelihood for internal promotion regardless of who is on the patent listing.

The next motivation to consider is emotional. The general question to consider here is whether not being listed on a patent dampens employee motivation to the degree that such a scheme is untenable. However, the specific question, one tractable only by considering the two-sided scale, is whether the dampening of motivation created by employee absence on a patent outweighs the decrease in market efficiency the current patent-induced inability to act on information

necessarily creates. We argue no. In the literature review[41] for their 2016 paper, *What Motivates Software Engineers?*[42], Sarah Beechem and John Hull identify the most commonly cited and strongest motivators for software engineers as ''the need to identify with the task' such as having clear goals, a personal interest, understanding the purpose of a task, how it fits in with the whole, having job satisfaction; and working on an identifiable piece of quality work." They additionally differentiate between the motivations of seasoned and newly-recruited engineers: "For example, the newly recruited Software Engineer could be highly motivated by job security and close supervision, whereas these same factors, especially close supervision, could turn out to be de-motivating to a seasoned Software Engineer. An experienced Software Engineer is more likely to be motivated by challenges, opportunities for recognition and autonomy." Of those motivations, only "opportunity for recognition" could be construed as a need to be listed on a patent. However, they did not observe this to be one of the most commonly cited motivations, and it would be a strange company indeed that could only provide opportunity for recognition through listing its employees on patents. There is a great deal of research that almost entirely identifies

[41]Sarah Beecham et al., *Motivation in Software Engineering: A systematic literature review*, 50 INFORMATION AND SOFTWARE TECHNOLOGY 860–878 (2008).

[42] Sarah Beecham & John Noll, *What Motivates Software Engineers Working in Global Software Development?*, PRODUCT-FOCUSED SOFTWARE PROCESS IMPROVEMENT LECTURE NOTES IN COMPUTER SCIENCE 193–209 (2015).

other incentives beyond patents and formal recognition as primary for software engineers motivation[43] [44] [45] [46].

## V) Implications of Sections II, III, and IV

Now is the time to superficially recap what we've laid out in the preceding sections. In Section II, we lay out a matching framework that claims if, on proper metrics, listing neither humans nor AI on the patents for AI-generated inventions is admissible, then either the patents should be granted to no one or the patent system is in need of reform. Section III argues that under existing law listing either AI or their creators on the AI-invented patents is not admissible. Section IV argues the same under a set of normative considerations. We feel these justifications make it abundantly clear that neither listing should be considered admissible under the current system. Our conclusion, however, is not to do away with patents entirely—as mentioned in Section II, we are amendable to a different underlying ethical justification of the patent system under which the elimination of patents is the optimal choice, however, this justification is neither our, nor the field's, accepted justification—but mend the patent system in a way more inline with the coming wave of AI inventions.

[43] Yunwen Ye & K. Kishida, *Toward an understanding of the motivation of open source software developers*, 25TH INTERNATIONAL CONFERENCE ON SOFTWARE ENGINEERING, 2003. PROCEEDINGS. (2003).

[45] J.m. Verner et al., *Factors that motivate software engineering teams: A four country empirical study*, 92 JOURNAL OF SYSTEMS AND SOFTWARE 115–127 (2014).

[46] Liva Teinberga & Darja Mite, *Towards Understanding of Software Engineer Motivation in Globally Distributed Projects*, 2011 IEEE SIXTH INTERNATIONAL CONFERENCE ON GLOBAL SOFTWARE ENGINEERING WORKSHOP (2011).

Finally, in an additional section, we look at other mismatches between AI-generated works and current intellectual property systems. These include the timing mismatch of AI and the length of intellectual property systems, as well as the mismatch between proprietary systems and the need for openness as artificial intelligence integrates with culturally expected legal norms. In terms of timing, the 20 years of patent protection and 120 years of copyright protection constitute an eternity in AI. The laws were designed for industries and actors that moved at a comparatively glacial pace. Consider that deep learning, which is the basis for the entire field of cutting-edge AI, was not practical until a 2006 paper opened the door for quickly training neural nets. Since then, the field has moved in leaps and bounds, analogous to what would be many lifetimes in other industries. In fact, the basis for most modern neural nets (generative adversarial models) only emerged in 2014.

From the perspective of openness, societal acceptance of intellectual property systems may well hinge on a degree of accessibility different from the norm in modern intellectual property systems. Currently, in terms of transparency, an inventor does not have to reveal very much about the software code to secure a patent or copyright. This can be extremely problematic in the changes looming ahead for artificial intelligence fields where public safety, liability, or criminal justice may be concerned.

Then what to do with the problem of inventorship, and more importantly, patent rights in the AI-generated work scenario? The old well-used methods are inapplicable to AI-generated work, in large part, because they were designed for a world in which nothing resembling AI-generated work existed. This is not to suggest that innovation in AI should be unprotected. Rather, instead of hewing to the current patent system as a model for the AI-generated work

patent system, we may be better off taking a page from data protection in pharmaceutical development. In that context, a brand-name drug company receives 4-5 years of exclusivity in exchange for making its safety and efficacy clinical trial data available to potential competitors. There is a shorter period of protection, enforced through the context of regulatory approval, in exchange for openness.

Taking this route provides a number of benefits. First, it allows us to avoid the impossibility of assigning patent rights. Rather than disfigure both the square peg and the round hole, we are better off finding method of integration. Finally, an approach of this sort also is better tuned to the pace of competition in the relevant field. Currently, there is an enormous first-mover advantage in the artificial intelligence space, and this advantage is particularly true for creative computers. In short, the problem of incentivizing AI-generated work is solvable; it simply requires thinking outside the box.