

Considering the Law as an Evaluative Mechanism for Computational Creativity

Stephen McGregor¹

Abstract. This paper presents a novel approach to the evaluation of creativity motivated by an examination of the current state of intellectual property law. Starting with an assessment of a significant decision by the Supreme Court of the United States, the legal parameters for what is considered creative are established. The ramifications of the legal situation for the practice of computational creativity are briefly considered. The question of what legally qualifies as creative is then reversed, turning to a consideration of whether the law might be considered a suitable mechanism for evaluating the inherent creativity in certain computational systems. A proposition for a new evaluative technique emerges, based on the presumption of a dualist world view endemic to the law as it stands. The implications of the proposed approach are considered; formalisation and implementation are left for the future.

1 Introduction

The law has long been concerned with creativity, with notions of the ownership of information and ideas dating back to ancient times, and the modern regime of intellectual property beginning several centuries ago. Though intellectual property laws are often portrayed by legislators as protecting the rights of people working to implement ideas, creators have often approached the law ambivalently, or sometimes even antagonistically. On the other hand, regardless of its shortcomings, the law offers the potential for a societal level of evaluation not readily available elsewhere. The law has emerged out of the complex interaction between the public and its governments over the course of history, and so in some ways the values of society are encoded in the very strictures by which it is regulated. On this basis alone a consideration of how the law might go about considering the creativity of a productive system seems worthwhile, even if the particulars of legal manifestations are found to be unacceptable on a philosophical level.

This paper represents essentially a thought experiment, and the objects of this experiment are two diametrically oriented questions regarding the law and computational creativity. On the one hand, how should the law approach the origin and ownership of artefacts which are generated by non-human agents? And, on the other hand, is the legal approach to authorship and creativity a potentially useful tool for addressing the hard issues surrounding the evaluation of computer generated artefacts? In other words, could the legal acknowledgement of the intellectual affiliation between a machine and its output be regarded as a sound criterion for the acceptance of a computational process as truly creative?

In what follows, these two questions will be explored one at a time, though the examination of the first will become the premise

for the response to the second. The domain of this exploration will, in particular, be literary works, and, correspondingly, copyright protection, though there are many reasons to believe a similar line of thought would apply to other types of creative output. The answer to the first question will ultimately have to be left to law makers and legal scholars to hash out, but it is a question which in the very asking raises some interesting issues with the law in its approach to the philosophy of mind. Pursuing the second question arguably leads to some valuable revelations about the nature of computational creativity itself, and it is here that the rewards of the proposed philosophical exercise will hopefully be discovered.

2 Ownership of Creative Artefacts

In 1983, an American business called Feist Publications attempted to buy a license from another business called Rural Telephone Service Company in order to include the directory listings provided for Rural's customers in a region of Kansas in a forthcoming compendium of listings covering a larger area. Rural rejected Feist's offer, but Feist used the listings in their publication anyway without Rural's consent. Rural subsequently sued Feist for copyright infringement, and the resultant litigation eventually percolated up through the American legal system until, by the beginning of 1991, it had reached the Supreme Court, the ultimate judicial body in the United States [1].²

The Supreme Court determined that Feist had not, in fact, unlawfully appropriated Rural's intellectual property, by way of the logic that the primary source compiled by Rural was effectively a list of facts, and lists of facts are fundamentally not original because they do not "possess at least some minimal degree of creativity," [1]. This ruling represented a decisive shift in the approach to intellectual property in American law, which had hitherto been predicated on the relatively simple, if flawed, principle of "sweat of the brow", meaning that the ownership of an informational artefact was somehow correlated with the degree of effort involved in the production of the artefact in question [1]. Under the new regime, it has become less clear where to draw the line between creative expressions, which are the intended targets of copyright protection, and concepts, which clearly can be equally creative but are crucially not intended as the target of copyright, and are only subject to patent law when they are realised in some practical application. Hence, following the Feist decision, the previously straightforward mechanisms for determining the ownership of informational artefacts has been thrown into confusion [10].

² This paper will focus on intellectual property law in the United States, in large part because a preponderance of literature on the situation in that country is available, but these findings will hopefully be construed as having somewhat universal applicability, particularly as international copyright law is to a large extent standardised under the Berne Convention.

¹ Queen Mary University of London, email: s.e.mcgregor@qmul.ac.uk

This vagueness of applicability has arguably opened the door to a provocative question which was not considered by the authors of the Feist decision: if creativity and originality are to be taken as the principle criteria for assigning copyright, how does this law apply when the author of the work in question behaved along these lines and yet is not a human but rather a computer? And this may not be as far fetched a question as it at first seems. The creativity criterion referenced above is almost immediately qualified in the Supreme Court decision: “To be sure, the requisite level of creativity is extremely low; even a slight amount will suffice. The vast majority of works make the grade quite easily, as they possess some creative spark.” [1]. Hence in defining the essential characteristic of creativity, even the Supreme Court has resorted to the indeterminate language of “sparks” standing in metaphorically for an undefined but presumably mental process.

The letter of the law evidently offers no guidance on how to identify the “creative spark” which makes a resultant artefact the property of its author [11]. What’s at work here seems to be a kind of type-token distinction, but not of Quine’s variety, where a type is “the class of all tokens,” [14]. Rather, the token here is a material outcropping of an entity rooted in an immaterial world, and the entity as a whole remains in a critical regard the property of the agent responsible for the generation of the type. Thus a text which someone physically possesses is merely a concrete manifestation still somehow tethered to something which exists in a unitary way in the same abstract mental space where the author’s “creative spark” resides: it is the type which is owned, and the token which is conditionally leased out.

What emerges from this analysis is a dualism fundamental to the prevalent state of intellectual property law. Creativity happens in a disembodied space, and only conceptual expressions which emerge from this space can be considered as belonging to a creator. When this approach is taken in light of the work of certain post-Cartesian philosophers [9, 13, 6] over the past several decades, it becomes abundantly clear that the Supreme Court has yet to come to terms with the recent history of the philosophy of language. To wit, the distinction between concept and expression which serves as the foundation of copyright [1] seems aggravatingly flippant to thinkers informed by the moves away from representational entanglement which have characterised many recent insights into both language and mind. So, to answer the question posed above, in order for a machine to be considered the author and owner of intellectual property, to the extent that this premise makes any sense at all, the machine would have to be considered autonomously creative. And in order to be considered creative, the machine would have to exhibit in some convincing way the type of mind/matter dualism which, at least in some circles, is so readily ascribed to human beings.

2.1 The Law is Discouraging CC

The concept of intellectual property is often couched in terms of motivating human creativity by protecting the livelihood of creators [1]. There have, however, been some eloquent expressions of dissent from this purported objective. Good arguments have been made, for instance, that innovation happens on a cooperative communal level [12], and that copyright laws restrict creators from using protected material in ways which would be broadly beneficial to society [17]. Regardless of whether the law helps or hinders human creators, though, there is evidence that the law discourages the pursuit of computational creativity.

In 2003, the multinational car manufacturer Toyota indirectly hired a design firm called Meshwerks to build digital models for

promoting a forthcoming line of vehicles. Meshwerks performed the task of building these models by basically wrapping the cars in a grid of tape and then mapping the points where the tape intersected across the surface of the vehicles onto a digital representation. The output of this procedure was a sort of pointillist 3D virtual sculpture of the surface of the car, which customers would be able to manipulate on a digital platform. The agreement between Toyota and Meshwerks allowed the models to be used on a one-off basis, but Toyota disregarded this deal and decided to use the models repeatedly in various marketing campaigns. Meshwerks sued Toyota for copyright infringement, but the lawsuit was ultimately rejected by the Tenth Circuit Court of Appeals in 2008 on the basis of the logic that the work done by Meshwerks did not meet the creativity criterion for intellectual property, citing the Feist decision as precedent in the matter [2].

Putting aside objections to the nature of copyright law in general, Toyota arguable did have a strong case here. The work done by Meshwerks was fundamentally not original; they were simply rendering a thing designed by Toyota in a different format. However, what’s interesting about this case from the perspective of computational creativity is not that Meshwerks’ output was considered uncreative, but the way in which Meshwerks attempted to defend their ownership of the models. Specifically, Meshwerks claimed that, while their initial process was cleverly mechanised, the output of that process was inadequate, and the bulk of the design work had to be done “by hand” [2].

So, in other words, the creativity stipulations of the canonical Feist decision appear to be engendering a situation where there is a motivation for creators to deny the role of computational agents in their processes. This is a portentous development for those interested in discussing ways in which computers might be construed as instantiating creativity, as it will only dissuade practitioners from acceding to such claims.

3 Creativity Judged by the Law

Of course in practice, legislators are generally not motivated by philosophical considerations one way or the other—indeed, copyright law seems to be more often than not an instrument for market manipulation and in particular a way for nations to leverage what is viewed as a competitive advantage in the information industry [3]. Furthermore, to the extent that these problems are considered on a theoretical basis, the idea of non-human authorship remains a fringe issue in the legal world, and, perhaps not surprisingly, to the degree that it has been addressed, the approaches proposed can be summarised as basically seeking to deny a machine’s capacity for creativity on the basis of its non-humanness. For instance, one author exploring this line of thought proposes an “independence requirement” which would seek to segregate the human effort in a situation where a human and a machine were working jointly to produce an artefact: the creativity of the systems output would be judged solely on the human contribution [11].

But putting aside arguments in support of the creative legitimacy of the computer’s work, the phrasing of this requirement is itself prescient: “the independence requirement focuses on the process of what the person in fact did instead of the end product,” [11]. Thus the very idea of judging humans differently than machines falls back on an examination of the process by which each agent acts in the world. In this regard, the legal move to analyse creativity in terms of its mode of production immediately calls to mind Boden’s own approach to the evaluation of non-human output, which focuses on the methodol-

ogy of traversing the conceptual space in which the creative artefact will be discovered rather than on the artefact itself [5]. In so doing, this legal proposal for an assessment of independence anticipates the second and more fundamental question explored in the present paper: can the legal framework for dealing with intellectual property inform the evaluation of computationally creative systems?

The law is in itself clearly an evaluative mechanism, designed both to guide the subjects of its strictures in what is permitted and to prescribe consequences when those strictures are violated. The importance of accommodating subjectivity in the evaluation of creative artefacts has been spelled out by Ritchie [15], but, on the other hand, evaluation by select committee can result in outcomes which are incomplete or biased. To the degree that law is designed to represent the values of the society to which it pertains, the judgement the law passes is reflective of a kind of consensus, and thus legal approval can potentially rise above the unsatisfactory arbitrariness of decisions made by small or unrepresentative groups. While such consensus should not be confused with universal acceptance, it can at least provide a valid point of reference in an individual's ultimate consideration of the value of an artefact in the context of that individual's community. An evaluation on a societal scale, for instance, poses an interesting counterpoint to evaluation by polling a sample group in the manner of a psychological study.

Wiggins has described computational creativity in terms of "behaviour exhibited by natural and artificial systems, which would be deemed creative if exhibited by humans," [19]. This should not be taken to imply a sort of "Turing Test", however, as "behaviour" here specifically refers to the methodology by which an artefact is generated, not to the sheer output—"behaviour" must not be confused with "behaviourism". Elsewhere, Wiggins has formalised Boden's approach to the field in a framework which includes a crucial mechanism for evaluation built into the creative system itself [18]. But, again per Ritchie's own formalism, the system's evaluative procedure is performed on its own operation in the process of creation [15]: the system is evaluating its own mechanism of production, and one of the key elements of that mechanism is the evaluation itself. As the evaluative mechanism becomes critically concerned with its own evaluation, something of a recursive crisis arises. Compounding this problem is the fact that, while what is desirable is an observation of methodology, it is unclear how the object of any assessment of a computational system could ever be performed on anything other than the system's external output.

The legal approach outlined in the previous section of this paper, for all its disheartening dualism, may actually offer some help in excavating creative systems from this quandary. To wit, the law suggests that an artefact is to be considered a creative product when its production has involved "some creative spark", establishing the precedence of a mind/matter divide in the legal definition of creativity. To extrapolate this criterion to a more general application, it would seem that the index of a creative system is the appearance of an agent which somehow simultaneously exists in two incompatible worlds, one mental and one physical. This move affords something which is, at least according to the principle of the law, an observable mechanism for evaluating creativity. The chief proposition of this paper is therefore this:

Proposition *The appearance of an evident dualism in a system, by which the system appears, upon analysis, to have internal states associated with its outputs, should be taken as a positive indicator that the system behaves creatively.*

No strong claim is made here about what would constitute evi-

dence of such a situation. Furthermore, it is hopefully quite clear that what is being suggested is not an acquiescence to a dualist perspective for the sake of defining creativity. To the contrary, one of the fundamental motivations in pursuing a positive case for computational creativity is to establish an empirical platform from which to attack the mystification of the mind which is characteristic of dualist philosophy, with its reliance on the material ineffability of mental representations. But the hope is that it is this very removal from representationalist theories of mind that will give the post-Cartesian thinker the leeway to employ a search for evident dualism as an evaluative mechanism without needing to take such appearances at face value.

In fact, the approach proposed here gets down to the hard question of why acceptance of a dualist world view has seemingly come so easily in the history of Western philosophy, a state of affairs which runs so deep that it is enshrined in the very letter of the law. It also arguably takes a step towards resolving the open question of defining what, to revisit Wiggins' terminology, counts as behaviour "which would be deemed creative if exhibited by humans", an issue which remains one of the most intransigent problems in the field. One of the confounding issues in computational creativity is the ease with which observers accept the evident mystery of human creativity, and conversely the understandable difficulty with which computers have in gaining this kind of acceptance. Through the identification felt by a human observer, a human creator is bestowed with the presumption of having feelings and intentionality, and the mere proclamation of phenomenology employed by some contemporary computational systems [8] is not a satisfactory solution here.

3.1 Proposal for Future Development

The construction of a framework for implementing this proposal for a test which measures the apparent dualism in a creative system is, by and large, left for development at a later date. As such, this paper is intended to serve simply as an indicator of a possible direction for exploration in the field. Nonetheless, a very rough suggestion for one way in which an advance on this prospective mechanism might be staged will be very briefly mooted here.

Previous approaches to the problem of evaluation have generally proposed either an exposure of the creative processes underwriting the computational system [7] or an acceptance that in the end all aspects of a creative system must ultimately be considered basically external output [16]. The approach proposed here must, similarly to the case for a revelation of the creative process, ultimately endorse a two-pronged assessment of what the system is doing. On the one hand, there must clearly be some sort of evaluation of the artefact itself, and this assessment would presumably look for the familiar properties of value, novelty, and surprise established as criteria for creative output [5]. Under the nascent approach being developed here, however, the second aspect of evaluation, the one which assesses the actual mechanism of production, would require a more active role for the observer: what is necessary is a kind of probing analysis of what kind of a system the creative agent really is, looking beyond its own explanation of its activity.

How might it be possible to investigate the nature of the operation of a symbol manipulating computational system, beneath the evident layer of the operational intentions with which it has been designed? What is called for seems to be a kind of metaphoric brain scan of a computer. In particular, it would be desirable to somehow look at a computational system and conclude that it seemed to be the kind of system which might maintain internal representations of its envi-

ronment. As it is not reasonable to posit some form of introspection for machines which have been specifically designed to manipulate symbols in a transparent and observer relative way, a different mechanism than directly enquiring the system about why it does what it does needs to be established. A framework for analysing the actual architecture of the system and of assessing this architecture *in situ* must be put forward.

Such a mechanism could presumably take on an uncountable number of different forms, and there is no reason to insist or even suspect that it would necessarily have to look for something in the creative agent that explicitly resembled a human brain or any other sort of biological entity; it would simply be searching for evidence of structures that have the appearance of internal representations. Having said this, one salient example of contemporary systems which seems to have the potential to meet the criteria for this kind of investigation are deep architectures. As the latest manifestation of the ongoing research into neural networks which has come in a number of waves since the mid-20th Century, deep architectures are specifically inspired by the highly distributed, decentralised, integrated nature of human brains.

Moreover, deep architectures are based on the premise that an effective system for interpreting data from an environment in semantically meaningful ways must have a large number of intermediate layers between the raw percepts and the high level perception of things as entities generally associated with consciousness [4]. Deep architectures establish a number of “factors of variation” on each compositional level of their model of the perceived world, and the various structures which exist on one representational level of abstraction can be combined in exponentially numerous ways to inform the next higher level of the architecture.

In this combinatory aspect of their operation, deep architectures have an intriguing correspondence with the exploratory approach to creativity laid out by Boden [5]. And, in their ability to continuously recalibrate the relationships between the representational spaces on each level, there may be the potential for higher order transformational creativity. Meanwhile, the very fact that such a system operates by passing representations of increasingly abstract degrees from one level to the next seems to fulfil the call for the semblance of dualism which this paper proposes should be found in a creative system. While the schema of deep architecture systems are, by nature, extremely complex, there should be a way to develop a methodology for exploring the system in its various stages of observation and identifying if it is building these types of representational structures, even if it is not feasible to know exactly what one of these structures specifically represents.

4 Conclusion

Going back to the source material, the problem of accepting the creativity of a computational system is further problematised in the language of the Feist decision itself, where facts are described as things which are “discovered” [1], and as such are not considered to be fundamentally creative. This comes uncomfortably close to Boden’s terminology regarding the search spaces of computationally creative agents and the exploratory aspect of the creative traversal of these spaces [5]—with the presumed objective of exploration being discovery.

This paper has proposed that a solution to the dilemma of persuading sceptical observers of the creative validity of computationally produced artefacts may be found in presenting systems which evince the kind of mind/matter dualism familiar from the Cartesian philosophical tradition. This approach has been motivated by an analysis

of prevalent attitudes towards copyright law, as embodied in influential decisions taken by high level American courts. These legal formulations are taken as representative of the prevalent state of affairs in the way society thinks about intellectual property.

It should be clear that the philosophical approach here is itself thoroughly anti-representational, and that in fact one of the primary motivations for pursuing computational creativity is to use it as a platform for investigating non-dualist angles on philosophy of mind. In fact, commitment to an anti-representational stance will guarantee that the impression of dualism will remain just that, a superficial impression. To a certain degree, this proposition falls back on the question, raised by Wiggins’ definition of creativity [19], of what exactly it means for a system to be somehow humanlike, but at the same time, by borrowing a dubious criterion from the public sphere of the law, this mechanism allows the theorist to step away from the issues surrounding the problem of who’s making judgements about a creative system without conceding to the suppositions inherent in those judgements.

At the same time, in dwelling on the potential value of false impressions for offering insight into computational creativity, there exists the potential for absolution from the crisis of externalism inherent to the evaluation of systems which are so fundamentally different from humans. *Pace* Colton et al, it is not sufficient, and probably not even appropriate, to try to assign phenomenology to computers by having them generate the kind of statements which humans make when they describe their creative motivations: such an approach admits to exactly the same problems as dualism itself. But by acknowledging the shortcomings of dualism, and by accepting the impossibility of ever *truly* talking about computations as anything other than observer-relative output, the double falsehood of internal mental states and computers with intentions and emotions resolve one another. The hope is that, from this resolution, a new mechanism for evaluation of computational creativity might emerge.

ACKNOWLEDGEMENTS

The research behind this paper was funded by EPSRC grant EP/L50483X/1. Thanks are also due to two reviewers for their insightful comments.

REFERENCES

- [1] Feist Publications Inc. v. Rural Tel. Service Co., 499 U.S. 330, 1991.
- [2] Meshwerks, Inc. v. Toyota Motor Sales U.S.A., Inc., 528 F.3d 1258 (10th Cir.), 2008.
- [3] Jonathan Band, ‘Armageddon on the potomac: The collections of information antipiracy act’, *D-Lib Magazine*, **5**(1), (1999).
- [4] Yoshue Bengio, ‘Learning deep architecture for ai’, *Machine Learning*, **2**(1), 1–127, (2009).
- [5] Margaret A. Boden, *The Creative Mind: Myths and Mechanisms*, Weidenfeld and Nicolson, London, 1990.
- [6] Robert B. Brandom, *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Harvard University Press, Cambridge, MA, 1994.
- [7] Simon Colton, ‘Creativity versus the perception of creativity in computational systems’, in *Proceedings of the AAAI Spring Symposium on Creative Intelligent Systems*, (2008).
- [8] Simon Colton, Jacob Goodwin, and Tony Veale, ‘Full-face poetry generation’, *Proceedings of the Third International Conference on Computational Creativity*, 95–102, (2012).
- [9] Donald Davidson, ‘Truth and meaning’, in *The Philosophy of Language*, ed., A. P. Martinich, 114–124, (2001).
- [10] Dennis S. Karjala, ‘Copyright and creativity’, *UCLA Entertainment Law Review*, **15**, (2008).
- [11] Edward Lee, ‘Digital originality’, *Vanderbilt Journal of Entertainment and Technology Law*, **14**(4), 919–957, (2014).

- [12] Lawrence Lessig, *The Future of Ideas: The Fate of the Commons in a Connected World*, Random House, New York, 2001.
- [13] Hilary Putnam, 'The meaning of "meaning"', in *The Twin Earth Chronicles: Twenty Years of Reflections on Hilary Putnam's "The Meaning of 'Meaning'"*, eds., Andrew Pessin and Sanford Goldberg, 3–52, M.E. Sharpe, Armonk, NY, (1996).
- [14] W. V. Quine, *Quiddities: An Intermittently Philosophical Dictionary*, The Belknap Press of Harvard University Press, Cambridge, MA, 1987.
- [15] Graeme Ritchie, 'Assessing creativity', in *Proceedings of the AISB Symposium on AI and Creativity in Arts and Science*, York, (2001).
- [16] Graeme Ritchie, 'Some empirical criteria for attributing creativity to a computer program', *Minds and Machines*, **17**(1), 67–99, (2007).
- [17] Siva Vaidhyanathan, *Copyrights and Copywrongs: The Rise of Intellectual Property and How It Threatens Creativity*, New York University Press, 2001.
- [18] Geraint A. Wiggins, 'A preliminary framework for description, analysis and comparison of creative systems', *Knowledge-Based Systems*, **19**, 449–458, (2006).
- [19] Geraint A. Wiggins, 'Searching for computational creativity', *New Generation Computing*, **24**, 209–222, (2006).