Article 25fa pilot End User Agreement

This publication is distributed under the terms of Article 25fa of the Dutch Copyright Act (Auteurswet) with explicit consent by the author. Dutch law entitles the maker of a short scientific work funded either wholly or partially by Dutch public funds to make that work publicly available for no consideration following a reasonable period of time after the work was first published, provided that clear reference is made to the source of the first publication of the work.

This publication is distributed under The Association of Universities in the Netherlands (VSNU) 'Article 25fa implementation' pilot project. In this pilot research outputs of researchers employed by Dutch Universities that comply with the legal requirements of Article 25fa of the Dutch Copyright Act are distributed online and free of cost or other barriers in institutional repositories. Research outputs are distributed six months after their first online publication in the original published version and with proper attribution to the source of the original publication.

You are permitted to download and use the publication for personal purposes. All rights remain with the author(s) and/or copyrights owner(s) of this work. Any use of the publication other than authorised under this licence or copyright law is prohibited.

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please contact the Library through email: copyright@ubn.ru.nl, or send a letter to:

University Library
Radboud University
Copyright Information Point
PO Box 9100
6500 HA Nijmegen

You will be contacted as soon as possible.

Mireille Hildebrandt*

LAW AS COMPUTATION IN THE ERA OF
ARTIFICIAL LEGAL INTELLIGENCE:
SPEAKING LAW TO THE POWER OF
STATISTICS†

The idea of artificial legal intelligence stems from a previous wave of artificial intelligence, then called jurimetrics. It was based on an algorithmic understanding of law, celebrating logic as the sole ingredient for proper legal argumentation. However, as Oliver Wendell Holmes has noted, the life of the law is experience rather than merely logic. Machine learning, which determines the current wave of artificial intelligence, is built on data-driven machine experience. The resulting artificial legal intelligence may be far more successful in terms of predicting the content of positive law. In this article, I discuss the assumptions of law and the Rule of Law and confront them with those of computational systems. As a twin article to my Chorley lecture on law as information, this should inform the extent to which artificial legal intelligence provides for responsible innovation in legal decision making.

Keywords: cybernetics, information theory, legal intelligence, legal protection by design, legal services, legal theory, meaning, political economy, Rule of Law, speaking law to power

1 *Introduction: does artificial legal intelligence challenge
the conceptual foundations of the law?*

Artificial legal intelligence has recently been described as a disruption of the legal services market, the legal profession, and prevalent business models, although scepticisms regarding overly broad claims and objections against the premature replacement of human legal expertise have surfaced.¹ Disruption of the legal profession, however, has been going on for

* Research Professor, Vrije Universiteit Brussel, Faculty of Law and Criminology, Brussels, Belgium; Full Professor, Radboud University, Science Faculty, Nijmegen, Netherlands. <http://orcid.org/0000-0003-4558-9149>

† This article was prepared for, and presented at, the Conference on Artificial Intelligence, Technology and the Future of Law at the Centre for Innovation, Law and Policy, University of Toronto, upon the invitation of Simon Stern, and first presented at the Innovation Policy Colloquium at New York University, upon the invitation of Helen Nissenbaum and Katherine Strandburg. Both presentations took place in March 2016, generating a host of salient comments that informed the final version.

1 Mainly positive: Richard Susskind, *The End of Lawyers? Rethinking the Nature of Legal Services*, rev ed (Oxford: Oxford University Press, 2010) [Susskind, *End of Lawyers*]; Daniel

quite some time, based on the commodification of legal services that have resulted from the practice of billable hours, combined with the complexity of an increasingly global business practice that finds itself confronted with overlapping jurisdictions, including major challenges in terms of dynamic tax law, environmental law, and more.² It seems that in the context of a transnational corporation it has become next to impossible to map the multi-level framework of partially contradictory legal requirements that determine corporate liability. At the same time, governments are struggling to enact a regulatory system that stays ahead of the myriad of legal constructions developed by corporate law firms to avoid taxation and other costly compliance issues.³ In addition, in some jurisdictions, access to legal counsel for those not extremely rich has become highly problematic due to the high costs of effective legal advice and representation.⁴

Thinking of law in terms of a market of legal services aligns with an implicit adherence to the basic tenets of law and economics. The latter goes with an understanding of law as just one particular – and maybe even cumbersome – example of governmental regulation that should prove its worth in comparison to other types of regulation that may be more effective and/or efficient.⁵ The concomitant ‘regulatory paradigm’

Martin Katz, ‘Quantitative Legal Prediction – or – How I Learned to Stop Worrying and Start Preparing for the Data Driven Future of the Legal Services Industry’ (2013) 62 Emory LJ 909 [Katz, ‘Quantitative Legal Prediction’]; Lawrence B Solum, ‘Artificial Meaning’ (2014) 89:1 Wash L Rev 69 [Solum, ‘Artificial Meaning’]; Anthony J Casey & Anthony Niblett, ‘Self-Driving Laws’ (2016) 66:4 UTLJ 429; Albert H Yoon, ‘The Post-Modern Lawyer: Technology and the Democratization of Legal Representation’ (2016) 66:4 UTLJ 456 [Yoon, ‘Post-Modern Lawyer’]; Benjamin Alarie, ‘The Path of the Law: Towards Legal Singularity’ (2016) 66:4 UTLJ 443 [Alarie, ‘Path of the Law’]. More or less cautious: Karen Yeung, ‘Algorithmic Regulation: A Critical Interrogation’ Regulation and Governance (2017), online: SSRN <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2972505>; Harry Surden, ‘Machine Learning and Law’ (2014) 89:1 Wash L Rev 87; Brian Sheppard, ‘Incomplete Innovation and the Premature Disruption of Legal Services’ (2015) 2015 Michigan State L Rev 1797; Dana Remus & Frank S Levy, ‘Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law,’ SSRN (27 November 2016), online: SSRN <<http://dx.doi.org/10.2139/ssrn.2701092>>.

² Yoon, ‘Post-Modern Lawyer,’ supra note 1 at 458–60.

³ ‘It Knows Their Methods: Watson and Financial Regulation,’ *The Economist* (22 October 2016), online: Economist Group <<https://www.economist.com/news/finance-and-economics/21709040-new-banking-rules-baffle-humans-can-machines-do-better-it-knows-their-methods>>.

⁴ Some of the rhetoric around new business models for legal services promises a democratization of legal knowledge – e.g., Yoon, ‘Post-Modern Lawyer,’ supra note 1; Suskind, *End of Lawyers*, supra note 1. If artificial legal intelligence is going to be part of the trade secret or intellectual property rights of big technology firms, this is not a likely scenario.

⁵ Still a canonical text: Richard A Posner, ‘The Decline of Law as an Autonomous Discipline: 1962–1987’ (1987) 100:4 Harv L Rev 761.

thinks in terms of influencing the behaviour of people,⁶ for which one could also use economic measures, techno-regulation, or various types of nudging, which may all prove more effective than legislation or adjudication.⁷ Such an economic understanding of law has been a turning point for the idea of lawyers as dignified stewards of individual justice and societal order. The idea that such stewardship has gone awry was aptly described in Anthony Kronman's *The Lost Lawyer*,⁸ although one can criticize its paternalistic and somewhat naive understanding of the role of lawyers and law firms in the nineteenth and twentieth centuries.⁹ Nevertheless, Kronman and others foresaw, or simply documented, the turn to commodification, which enables legal scholars to discuss 'the law' as being equivalent with 'the legal services market,' thus validating the Chicago schools of law and economics (whether based on rational choice and game theory or on the new school of behavioural economics).

If the perspective is economics, however, I would prefer an analysis of the political economy that drives the 'legal services market,' because I am not sure we should take the assumptions underlying the narrative of commodified legal services for granted.¹⁰ Instead of investigating how artificial intelligence will reconfigure the legal services market, I would like to ask the preliminary question of how law does and should constitute and regulate such economic markets, while sustaining the checks and balances of the Rule of Law that should define them.¹¹

6 Julia Black, 'Critical Reflections on Regulation' (2002) 27 Australian Journal of Legal Philosophy 47.

7 Lawrence Lessig, *Code Version 2.0* (New York: Basic Books, 2006); Ronald Leenes, 'Framing Techno-Regulation: An Exploration of State and Non-State Regulation by Technology' (2011) 5:2 *Legisprudence* 143; Cass R Sunstein, *The Ethics of Influence: Government in the Age of Behavioral Science* (Cambridge, UK: Cambridge University Press, 2016).

8 Anthony T Kronman, *The Lost Lawyer: Failing Ideals of the Legal Profession* (Cambridge, MA: Harvard University Press, 1995); Mary Ann Glendon, *A Nation under Lawyers* (Cambridge, MA: Harvard University Press, 1996); Frank Pasquale, 'Automating the Professions: Utopian Pipe Dream or Dystopian Nightmare?', *Los Angeles Review of Books* (15 March 2016), online: Los Angeles Review of Books <<https://lareviewofbooks.org/article/automating-the-professions-utopian-pipe-dream-or-dystopian-nightmare/>>.

9 E.g., from a feminist perspective, see Minna Kotkin, 'Professionalism, Gender and the Public Interest: The Advocacy of Protection' (1995) 8 *St Thomas L Rev* 157.

10 Julie E Cohen, 'The Regulatory State in the Information Age' (2016) 17:2 *Theoretical Inquiries in Law* 369. Dana Remus, 'Reconstructing Professionalism,' UNC Legal Studies Research Paper No. 2676094 (30 April 2015), online: SSRN <<https://ssrn.com/abstract=2676094>>.

11 I capitalize 'Rule of Law' not merely to pay my respects but, foremost, also to prevent confusion with a particular rule/norm of law. Cf Jeremy Waldron, 'The Concept and Rule of Law' (New York University School of Law Public Law Research Paper no. 08-50, 24 September 2008) at 1.

My research is focused on how law co-determines the playing field of economic markets and on how it informs the constitutional or constitutive relationship between the public, the social, and the private that is the condition of possibility for a durable, fair, and thriving market place. Instead of looking at law as a product or service, or as something that regulates a given population, I investigate law as constitutive of societal defaults, notably those defaults that ground an effective Rule of Law. The latter entails a system of checks and balances that gives effective standing to those subject to the law, affording them effective legal remedies to contest a violation of their fundamental rights.¹² Based on this perspective, I will provide a tentative answer to the question of whether artificial legal intelligence challenges the conceptual foundations of the law.

In Part II, I dig into the conceptual foundations of the law, acknowledging that these foundations co-determine what Bruno Latour has termed the ‘mode of existence’ of the law. The inquiry has been inspired by my Chorley lecture, which investigated ‘law as information in the era of data-driven agency.’¹³ It entails investigating the concept of information, distinguishing the mathematical theory of information that grounds both computer science and artificial intelligence from the notion of meaningful information that is aligned with human language, written and printed text, and modern positive law. Based on theories of human language usage, I discuss the mode of existence of law in terms of speech act theory, highlighting the performativity of positive law, which elucidates law as information in the sense of meaning generating.

In Part III, on law as information in the computational sense, I confront the assumptions of the mathematical theory of information that underpins artificial legal intelligence. This Part will highlight the concept of performance as being pivotal for machine intelligence and will examine how machine performance may affect law’s performativity. Finally, I will conclude by suggesting that the concept and practice of ‘legal protection by design’ may contribute to sustaining law as meaningful information that informs the consequences of our actions, seeking to bring artificial legal intelligence under the Rule of Law.

¹² Within Europe, the right to an effective remedy against violations of human rights is articulated in art 13 of the *Convention for the Protection of Human Rights and Fundamental Freedoms*, 4 November 1950, 213 UNTS 221 (entered into force 3 September 1953).

¹³ Mireille Hildebrandt, ‘Law as Information in the Era of Data-Driven Agency’ (2016) 79:1 *Modern L Rev* 1.

A INFORMATION, MEANING, AND PERFORMATIVITY

According to Raphael Capurro and Birger Hjørland, the concept of information comes from the latin '*in-formare*,' which refers to 'one thing having a formative impact on another thing.'¹⁴ When I say that my lifestyle informs my health I am not suggesting that my lifestyle provides information about my health but, rather, indicating that my lifestyle influences or even transforms my health. Precisely because my lifestyle impacts my health, it may also provide information about my health. Capurro and Hjørland trace this second meaning, which refers to 'communicating something about something else,' to more recent changes in the use of the notion of information. One could say that information either refers to 'formativeness' or 'aboutness,' while acknowledging their interrelationship. Data that is stored in a computing system without being used can be seen as information about something that, until it is employed in some way or another, has no formative impact as such. Nevertheless, it may have already informed the choices people make, or be the result of such choices, and it may still come to inform the algorithms that are trained on the data.

Before moving into the realm of computational systems, we need to ask how human beings share and digest information as well as how their actions are informed by the information they gather while navigating their world.¹⁵ This 'world' is not just a matter of bricks and bones but also consists, to a large extent, of concepts and institutions that inform the consequences of our actions. In the example of lifestyle informing my health, one could perhaps replace 'inform' with 'cause,' but this is not always possible. Often, the formative aspect of the relationship is based on the performative nature of specific types of language usage. For instance, money is not, or no longer, defined by the value of its physical embodiment (such as with gold coins in an earlier age). Since promissory notes turned into paper money, the concept of money depends on a shared understanding of what counts as money. In a sense, 'money' is

14 Rafael Capurro & Birger Hjørland, 'The Concept of Information' (2003) 37:1 Annual Review of Information Science and Technology 343.

15 Navigation is a pivotal aspect of intelligence for mobile beings, as it requires the anticipation of collusion, see e.g. Rodney Brooks, 'Intelligence without Reason' (Paper for Computers and Thought: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, Massachusetts Institute of Technology, 1991) at 569; Rodney Brooks, 'Intelligence without Representation' (1991) 47 Artificial Intelligence 139. 'Higher' intelligence derives from navigating our institutional environment.

built on thin air.¹⁶ More to the point, it is built on a particular usage of human language. As long as we all agree that the coin, the bill, or the numbers in one's electronic bank account count as a specific type of currency, these coins, bills, and electronic bank statements exist as money. This is not the result of a postmodern relativism that wrongly suggests that anything and everything is socially constructed (which is, of course, nonsense) but, rather, a consequence of the performative nature of language usage.¹⁷

John Austin and John Searle have highlighted the fact that speech or text does not merely describe but often does what it describes.¹⁸ They developed speech act theory, distinguishing between locutionary speech acts (which contain propositions – 'aboutness') and illocutionary speech acts (which perform what they refer to – 'formativeness'). When a civil servant pronounces the words 'I marry thee,' he is not describing what he is doing but, instead, actually doing it. The speech here is an act, it performs the act of marrying a couple, and this is just one example of the performative nature of human language use. When I write out a cheque to another person, I am not describing the transfer of money but actually instructing my bank to provide money to that person upon her request. The piece of paper counts as such an instruction and also counts as a right for that person to receive the money. The cheque does not merely contain information about this instruction, but it also transforms the piece of paper into a token that will inform the actions of the bank and the beneficiary.

Speech act theory compares well with another inquiry into the performative effects of human language usage – namely, semiotics or the study of how signs generate meaning.¹⁹ Semiotics demonstrates that signs never

16 Pierre Lévy, *Becoming Virtual: Reality in the Digital Age* (New York: Plenum Trade, 1998) at 68–70. An interesting question is whether Bitcoin depends on the use of language in the same way as 'traditional' money since it seems to entirely depend on self-executing code.

17 Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network-Theory*, Clarendon Lectures in Management Studies (Oxford: Oxford University Press, 2005); Bruno Latour, *An Inquiry into Modes of Existence: An Anthropology of the Moderns*, trans Catherine Porter (Cambridge, MA: Harvard University Press, 2013).

18 J.L. Austin, *How to Do Things with Words*, 2d ed (Cambridge, MA: Harvard University Press, 1975); John R Searle, *Speech Acts, an Essay in the Philosophy of Language* (Cambridge, UK: Cambridge University Press, 1969).

19 Doede Nauta, *The Meaning of Information* (Berlin: Mouton, 1972). Semiotics shares some insights with Peircean pragmatism, as Pierce was the founding father of both. See Charles Saunders Peirce, *Selected Writings*, ed by Philip P Wiener (New York: Dover, 1958). Note that the pragmatist maxim defines the expected consequences of the use of a concept as its meaning, highlighting the connection between expectancy theories of language use and its performative effect.

mean anything until they are used by interacting human beings. Language usage implies connecting a system of signs with whatever these signs stand for, refer to, or perform. The word 'cup' can be a sign that stands for, or refers to, a physical contraption that allows us to carry liquids, while also enabling us to drink the liquid. To understand the meaning of a cup, we need to be aware of other terms, such as liquid, drinking, glass, and so on. The word 'education' stands for, or refers to, not only buildings, books, teachers, and students but also a whole network of other terms. Both 'cup' and 'education' depend on shared habits, patterns of behaviour, and the anticipation that practices such as drinking from a cup and enrolling for an education will make sense to others who live in the same shared world.²⁰ One of the most interesting features of this shared world, made of bricks and bones, words and habits, and fine-tuned mutual expectations, is a persistent discontinuity, an iterant threat of disruption, and an abyss that opens up whenever expectations are violated or ignored.

Since the meaning of our behaviours, including our words, depend on the web of signs and what they refer to, stand for, or 'make up' (as with money), meaning is always shifting. Keeping it stable requires hard work. Meaning is ultimately a matter of joint performance (I cannot create meaning on my own).²¹ It is inherently unstable (always depending on how common sense develops) and, thereby, creates room for productive misunderstandings and different interpretations and reflections on what it is we mean with a word, a concept, a promise, a request, or even an order.²² Therefore, information as aboutness will not do; to sustain our world, information must be formative, and, when it comes to concepts and institutions, it must have illocutionary (performative) force. It is important to note that this particular type of formativeness (constituting meaning) ties in with *human* language use; we cannot assume that other organisms or data-driven artificial agents are capable of generating meaning in this particular sense.

B LAW AS INFORMATION: THE PERFORMATIVE NATURE OF POSITIVE LAW

1 *the force of law*

As should be clear, speech act theory is particularly relevant for the law. The performative aspects of language usage that inform the mode of

²⁰ See 'Welt and Umwelt' in Mireille Hildebrandt, *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology* (Cheltenham: Edward Elgar, 2015) 51, section 3.3.3.1.

²¹ Ludwig Wittgenstein, *Philosophical Investigations*, 4th rev ed (Malden, MA: Wiley-Blackwell, 2009).

²² Judith Butler, *Giving an Account of Oneself* (New York: Fordham University Press, 2005) [Butler, *Giving an Account*].

existence of institutions such as economic markets, marriage, or the transfer of property all depend on 'the force of law' or what lawyers call 'legal effect.' The institution of economic markets is enabled by legal precepts such as the freedom to contract and various types of private law liability. The legal effect of a marriage is far reaching, notably with regard to creditors and inheritance, and the transfer of property is defined by a whole series of legal conditions that must be met to actually differentiate between 'mere' possession and ownership. From a social science or economic perspective, one might believe that the force of law consists of enforcement, by means of court injunctions, finally depending on the power of police (brute force). This, however, is a fatal misunderstanding. The force of law depends on the interlocking of performative speech acts (which are often written, as with legal deeds, legislation, and judgments) and the monopoly of violence that looms in the background to ensure that legal effect has the kind of teeth that moral prescriptions lack. If law were only dependent on enforcement, it would not be law but, rather, administration, discipline, or simply a matter of violence. This relates to the idea of exclusionary rules that relieve people from having to develop their own reasons for compliance or non-compliance and connects with the fact that fear of enforcement cannot be the sole reason to follow legal rules. Even if it could be the sole reason to follow up on disciplinary, military, or administrative orders, legal rules assume a legitimacy that does not accord with brute force.²³

Legal effect is generated by legal norms, which must be derived from the sources of law, notably legislation, binding international treaties and case law (the formal sources of law), and legally relevant custom, fundamental legal principles, and doctrine (the informal sources of law) that help to interpret the legally binding texts of legislation, treaties, and case law. Clearly, both types of sources of law interact with, and inform, each other. The sources of law, in that sense, do not merely contain information about the law but also institute the law as such. Although legal knowledge can be seen as information about the legal effect of one's actions (implying a propositional logic and locutionary speech acts), to understand legal effect we must pay attention to how law informs (conditions) the consequences of our actions (implying a learning process and illocutionary speech acts).

²³ Obviously, law is in part instituted by violence, see e.g. Robert Cover et al, *Narrative, Violence, and the Law: The Essays of Robert Cover* (Ann Arbor: University of Michigan Press, 1995). Reducing it to violence, however, makes no sense. See e.g. Jacques Derida, 'Force of Law: The "Mystical Foundation of Authority"' (1990) 11 Cardozo L Rev 920. On exclusionary rules, see Joseph Raz, 'Reasons for Action, Decisions and Norms' (1975) 84:336 Mind 481.

2 *legal certainty, predictability, and openness*

We can trace a keen awareness of the informative and performative force of the law in Oliver Wendell Holmes's famous remarks that 'the life of the law has not been logic; it has been experience,'²⁴ defining law as 'the prophecies of what the courts will do in fact and nothing more pretentious.'²⁵ This links law with experience and prediction and, thus, with its role in reducing the uncertainty we may encounter about the consequences of our actions. Holmes may even be seen as a direct precursor of both law and economics and artificial legal intelligence, stating that 'for the rational study of the law the blackletter man may be the man of the present, but the man of the future is the man of statistics and the master of economics.'²⁶ As Ian Kerr has argued, however, Holmes's argument is about the opposition between logic (inexorable certainty) and statistics (calculated uncertainty), highlighting the crucial importance of providing those subject to law with a means to stabilize expectations, thus enabling people to go about their daily lives.²⁷ As a practising lawyer, Holmes knew that legal certainty is never given; it depends on the foreseeability of human action and human interpretation, both of which are fundamentally uncertain. Therefore, it is a matter of statistics rather than logic, of experience and learning rather than supposedly objective Framers' intention or literal textual interpretation. This implies that legal certainty is dynamic and requires hard work to be sustained. If this is done well, the uncertainty that grounds legal certainty is not a drawback but, rather, an asset; it makes sure that legal certainty does not trump justice and instrumentality, two other aims of the law that are at odds with a rigid interpretation of legal text.²⁸ Uncertainty is inherent in what Herbert Hart famously described as the open texture of fundamental legal concepts, safeguarding the creative and highly artificial nature

24 Oliver Wendell Holmes, *The Common Law* (Mineola, NY: Dover Publications, 1991) at 1. This ties in with the pragmatist maxim and an expectancy theory of legal norms. See JFG Glastra van Loon, 'Rules and Commands' (1958) 67:268 *Mind* 1.

25 Oliver Wendell Holmes, 'The Path of the Law' (1997) 110 *Harv L Rev* 994.

26 *Ibid* at 1001.

27 Ian Kerr, 'Chapter 4: Prediction, Preemption, Presumption: The Path of Law after the Computational Turn' in Mireille Hildebrandt & Katja de Vries, eds, *Privacy, Due Process and the Computational Turn: The Philosophy of Law Meets the Philosophy of Technology* (New York: Routledge, 2013) 91.

28 On the tension between these three 'antinomian' goals, see Gustav Radbruch, 'Legal Philosophy' in Kurt Wilk, ed and trans, *The Legal Philosophies of Lask, Radbruch and Dabin* (Cambridge, MA: Harvard University Press, 2014) 44; Mireille Hildebrandt, 'Radbruch's Rechtsstaat and Schmitt's Legal Order: Legalism, Legality, and the Institution of Law' (2015) 2:1 *Critical Analysis of Law* 42 [Hildebrandt, 'Radbruch's Rechtsstaat'].

of human intercourse as well as legal constructs.²⁹ As John Dewey has noted, the artificial nature of law must not be equated with its supposedly fictitious nature.³⁰ An artificial lake is not an imaginary lake but, instead, an act of creation and construction; the discussion should be on how to create and what to construct.³¹ This also involves the relationship between performativity and legal certainty.

In an interesting article in a previous issue of this journal, Benjamin Alarie emphasizes the potential of machine learning when it comes to ensuring legal certainty, taking an example from an overly complicated tax law, notably the distinction between an employee and an independent contractor.³² The distinction is highly relevant for the legal effect generated by tax law, both for corporations and for the individual persons that are either employed or contracted. The corpus of case law that helps to inform us about the distinction is so very extensive, casuistic, and complicated that many billable hours must be spent to advise corporations. Obviously, the space for negotiation on the side of the employee/contractor is not so great, and the consequences of misqualification as either an employee or an independent contractor may be hazardous for them also (think of non-declared profits, fines for not keeping administration). According to Alarie, artificial legal intelligence enables a fast and refined prediction of the relevant legal effect, based on a granular analysis of relevant case law by machine learning algorithms. Algorithms are trained on legal text (case law) by legal experts and data scientists, optimizing the mathematical relationship between input (the corpus of relevant legal texts) and output (the judgments) to a point where the algorithms get it right 98 per cent (which is better than the human lawyers who trained the algorithms). Note that the machines are merely simulating mathematically what human reasoning has come up with in written legal text.

Nevertheless, achieving legal certainty as to one's legal obligations under tax law is a crucial condition for individual freedom, as Jeremy Waldron has insisted:

There may be no getting away from legal constraint in the circumstances of modern life, but freedom is possible nevertheless if people know in advance how the law will operate and how they have to act if they are to avoid its

²⁹ HLA Hart, *The Concept of Law* (Oxford: Clarendon Press, 1994).

³⁰ John Dewey, 'The Historic Background of Corporate Legal Personality' (1926) 35:6 Yale LJ 655 at 655–6, n 1.

³¹ Ibid.

³² Alarie, 'Path of the Law,' *supra* note 1 at 448–51.

application. Knowing in advance how the law will operate enables one to make plans and work around its requirements.³³

However, whereas Alarie seems to believe that quasi-mathematical predictive accuracy is the final goal of legal advice, Waldron continues:

The institutionalized recognition of a distinctive set of norms may be an important feature. But at least as important is what we do in law with the norms that we identify. We don't just obey them or apply the sanctions that they ordain; we argue over them adversarially, we use our sense of what is at stake in their application to license a continual process of argument back and forth, and we engage in elaborate interpretive exercises about what it means to apply them faithfully as a system to the cases that come before us.³⁴

Waldron refers to the renowned legal philosopher Neil MacCormick as

pointing out that law is an argumentative discipline and no analytic theory of what law is and what distinguishes legal systems from other systems of governance can afford to ignore this aspect of our legal practice, and the distinctive role it plays in a legal system's treating ordinary citizens with respect as active centers of intelligence.³⁵

This response is no surprise as MacCormick applied speech act theory to better understand the nature of law and its crucial role in shaping our institutional environment.³⁶ Law as information, therefore, must be understood as referring to law as a coherent web of speech acts that inform the consequences of our actions, itself informed by the triple tenets of legal certainty, justice, and instrumentality that hold together jurisdiction (the force of law), community (even if between strangers), and instrumentality (the policy objectives of the democratic legislator). The antinomial tension between these three tenets require a trained and reflective judgment that is neither mathematically consistent nor based on subjectivist psychology; the judgment is defined by the performative effect it anticipates.

Alarie is betting on what he calls 'legal singularity,' the moment that artificial legal intelligence will get it right all of the time, thus eradicating

³³ Jeremy Waldron, 'The Rule of Law and the Importance of Procedure' (New York University Public Law and Legal Theory Working Paper, 2010), online: Nelco Legal Scholarship Repository <http://lsr.nelco.org/nyu_plltwp/234> at 19 [Waldron, 'Rule of Law'].

³⁴ Ibid at 20.

³⁵ Ibid.

³⁶ Neil MacCormick, *Institutions of Law: An Essay in Legal Theory* (Oxford: Oxford University Press, 2007).

any and all legal uncertainty.³⁷ He seems to mistake, however, the mathematical simulation of legal judgment for legal judgment itself. Whereas machines may become very good in such simulation, judgment itself is predicated on the contestability of any specific interpretation of legal certainty in the light of the integrity of the legal system, which goes way beyond a quasi-mathematical consistency.³⁸ Alarie refers to John Rawls's 'reflective equilibrium' as the final touchstone for legal singularity, which assumes (a) that simulation is equivalent with what it simulated and (b) that, while legal judgment may be intractable at this moment in time, it is computable and will become tractable in due course.³⁹ The second assumption ignores the fact that machine learning can be performed in multiple ways, leading to different outcomes. The fact that computability can be 'done' in different ways will 'simply' displace the adversarial debate to the level of training sets, types of algorithms, and the design of the hypotheses spaces.⁴⁰ Moreover, the very fact that computability implies translation and simulation also implies that judgments are not computable if computability is understood as a perfect translation. The first point (mistaking a simulation for what is simulated) turns on the difference between the performativity of a legal judgment by a court of law and the performance rating of machine-learning algorithms. These are two entirely different 'things,' requiring a more in-depth inquiry into law as computation.

III *Law as computation*

A COMPUTATION, FEEDBACK, AND PERFORMANCE

The mathematical theory of information that grounds computer science builds on the second meaning of information (communicating aboutness);

³⁷ This is a reference to Kurzweil's expectation that machine intelligence will supersede human intelligence, creating a hybrid intelligence where human and machine can no longer be distinguished. Cf Ray Kurzweil, *The Singularity Is Near: When Humans Transcend Biology* (New York: Viking, 2005).

³⁸ On the complexities of matching legal certainty with justice and purpose in Radbruch's legal philosophy, see Hildebrandt, 'Radbruch's Rechtsstaat,' supra note 28. On the difference between consistency and integrity, see Ronald Dworkin, *Law's Empire* (Glasgow: Fontana, 1991).

³⁹ Something is computable if it can be translated into machine-readable data. Due to complexity, problems that are computable may nevertheless be intractable when it would take too long to calculate a solution.

⁴⁰ Alarie seems to recognize this and foresees a 'legal arms race, with the machines playing a central role in the escalation of tensions' between governments and professional legal advisors working with machine-learning systems. Alarie, 'Path of the Law,' supra note 1 at 451.

it aims to describe how certain discrete bits of content (aboutness) can be transferred from A to B as fast as possible, at low cost, without changing the content, and without providing access to anyone other than either A or B. The mathematical theory of information was developed by Claude Shannon during World War II, with an eye to the fast and reliable transfer of secret messages.⁴¹ Shannon was not interested in the content itself but, rather, focused on a type of encryption that would ensure the sameness of the content with as few signals as possible (compression). To achieve the sparsest use of signals, he defined information in terms of surprise and uncertainty; his goal was to calculate which signals could be left out without distorting the content. In the end, this resulted in a mathematical definition of information, based on the concept of entropy.⁴² A data set with an entirely random distribution has complete entropy; the uncertainty is 100 per cent because it is not possible to make any prediction about the distribution. Any deviation of random distribution reduces the entropy and, thus, reduces uncertainty. To understand the potential impact of data-driven applications on the law, we should remember that computer science is built on a separation between signs and their meaning that is radically different from the semantics of human language.

The relationship between information and uncertainty was taken up by Norbert Wiener, the founding father of cybernetics, who was not interested in transferring secret information but, instead, in exercising control from a distance.⁴³ *Cyber* is Greek for ‘to steer’; cybernetics is the art or science of remote control over an environment. Wiener focused on how feedback from the environment allows an agent (whether an organism or a machine) to adapt its own behaviour to achieve its goals. The idea is that feedback reduces uncertainty about how the environment may respond to the agent’s behaviour; feedback is information about the past and the present that supposedly informs the future. Herbert Simon, one of the architects of the science of artificial intelligence, explored the crucial role of pattern recognition in feedback.⁴⁴ For an

41 Claude E Shannon & Warren Weaver, *The Mathematical Theory of Communication* (Urbana: University of Illinois Press, 1975).

42 Ibid.

43 Norbert Wiener, *The Human Use of Human Beings: Cybernetics and Society* (Cambridge, MA: Da Capo Press, 1988). Shannon and Wiener defined information in opposite ways; Shannon equated information with maximum surprise and, thus, with entropy (random distribution), whereas Wiener equated information with reduction of uncertainty and, thus, with neg-entropy (structured data). Common sense follows Wiener, who states that information is what reduces uncertainty; the greater the uncertainty, the greater the need for information.

44 Herbert A Simon, *The Sciences of the Artificial* (Cambridge, MA: MIT Press, 1996).

organism or a machine to successfully adapt its behaviour, it must anticipate how current regularities may play out in the future. Patterns are not the same as regularities; whereas regularities are about the past, patterns aim to reach into the future. They are speculative as they need to move beyond mechanical application.

This brings us to the sub-discipline of machine learning (ML): that is the use of computing systems to detect patterns in data that allow a system to update its own program. The idea is simple; to reach its goals, the system must reduce uncertainty about the effects of its own behaviour. To reduce such uncertainty, it probes its environment and processes the feedback, reconfiguring its processing algorithms until its goal is reached. Here we return to mathematics, though not to encryption. ML seeks to infer correct mathematical functions to describe patterns in a data set or between given input data (for instance, data on speed or brake behaviour) and given output data (for instance, machine readable 'safe driving behaviour' for an autonomous car). The result is a set of, potentially reconfigured, instructions that produces the desired behaviour of the system. To develop the right mathematical function, data scientists will construct a so-called hypotheses space, consisting of a number of mathematical functions (hypotheses about the relationship between data in a data set or between input and output) that are tested, tweaked, and tested again, until the system gets it right most of the time.⁴⁵ In his handbook on ML, Tom Mitchell writes: '[A] machine learns with respect to a particular task T, performance metric P, and type of experience E, if the system reliably improves its performance P at task T, following experience E.'⁴⁶

Much depends on who decides whether the system 'gets it right,' which in turn depends on the definition of task T (in machine-readable terms) and performance metric P (which determines progress). Obviously, the idea of performance has no relation with that of the illocutionary force or performativity of speech acts. The performance measure

⁴⁵ On the idea that good enough may be better than best, see Herbert Alexander Simon, *Models of Bounded Rationality: Empirically Grounded Economic Reason* (Cambridge, MA: MIT Press, 1982) at 295ff (discussing 'satisficing' as often superior to 'optimization'). In the same line, see Christoph Engel & Gerd Gigerenzer, 'Law and Heuristics: An Interdisciplinary Venture' in Gerd Gigerenzer & Christoph Engel, eds, *Heuristics and the Law* (Cambridge, MA: MIT Press, 2006) 1. On the Neoplatonic idea that mathematics discloses the underlying 'true' nature of reality and on how this has influenced modern science from the sixteenth century and, more recently, data science as a new – machinic – modality of Neoplatinism, see Dan McQuillan, 'Data Science as Machinic Neoplatonism' (2017) 30 *Philosophy and Technology* 1, online: Springer <<http://link.springer.com/10.1007/s13347-017-0273-3>>.

⁴⁶ Thomas Mitchell, *Machine Learning* (New York: McGraw-Hill Education, 1997) at 2.

is a machine-readable variable that enables the ranking of the machine's behaviour. Similarly, the reliability of ML depends on a number of inevitable trade-offs, notably between volume, relevance, completeness, accuracy, and correctness of the training data set, the dimensionality and aptitude of the hypotheses space, the time taken for iterant testing, and the availability of the relevant domain expertise (how do we know when the machine is 'getting it right?').

Some have suggested that information is data plus meaning.⁴⁷ This highlights the fact that a piece of datum does not speak for itself; it requires the attribution of meaning by those who use the data, once they use it. It also suggests that information is always meaningful in the above sense. This is clearly not the case. Although many organisms do not speak our type of language, they all depend on information to survive and flourish. Artificial intelligence similarly depends on the processing of information (in the form of digital data points). Neither in the case of organisms nor in the case of intelligent machines does information necessarily imply the attribution of meaning, as it may in the case of humans. Machines work with signs, they do not speak human language; they 'live' in a 'world' made of hardware and software, data, and code. Their perception is limited to machine-readable data, their cognition is based on computation and manipulation of signs, not on generating meaning (though we may attribute meaning to their output and/or their behaviour). The mode of existence of meaningful information in human society builds on the spoken word and on handwritten and printed text, each with its own affordances.⁴⁸ Meaning differs from information in its dependence on consciousness, self-consciousness, and unconsciousness (which depends on consciousness, by definition). Meaning thus depends on the curious entanglement of self-reflection, rational discourse, and emotional awareness that hinges on the opacity of our dynamic and largely inaccessible unconscious.⁴⁹

Data and code – and the hyperconnectivity they feed on – do not attribute meaning. Although ML has generated feedback mechanisms, these are not similar to human self-reflection.⁵⁰ Computer systems are now

⁴⁷ Luciano Floridi, *Information: A Very Short Introduction* (Oxford: Oxford University Press, 2010).

⁴⁸ Jack Goody, *The Logic of Writing and the Organization of Society* (Cambridge, UK: Cambridge University Press, 1986); Walter Ong, *Orality and Literacy: The Technologizing of the Word* (London: Methuen, 1982); Elisabeth Eisenstein, *The Printing Revolution in Early Modern Europe* (Cambridge, UK: Cambridge University Press, 2005).

⁴⁹ Butler, *Giving an Account*, supra note 22. Siri Hustvedt, *A Woman Looking at Men Looking at Women: Essays on Art, Sex, and the Mind* (New York: Simon & Schuster, 2016).

⁵⁰ Alarie, 'The Path of the Law,' supra note 1 at 454, speaks of an artificial legal intelligence that achieves Rawls's 'reflective equilibrium.' It can, however, only attempt to

capable of second order – if not n-order – reconfigurations, but they have no self-conscious awareness, they cannot suffer pain or humiliation, even where they excel at simulating such emotions.⁵¹ We are now living with creatures of our own making that can anticipate our behaviours and pre-empt our intent. They inform our actions even if we are not aware of this, and their inner workings seem as opaque as our own unconscious. We have little information about their operations, and it is our (largely unconsciously automated) behaviours that inform their often-inscrutable decisions.⁵² Computing systems perform certain functions, based on computation. Their decisions have a normative effect in the sense that they will induce, enforce, inhibit, or rule out some of our actions. The crucial issue will be how to differentiate their normative force from the force of law and how to make sure we do not confuse the performativity of positive law with the performance of artificial legal intelligence.⁵³

B LAW AS COMPUTATION: DATA-DRIVEN LAW

1 *challenging law and the rule of law*

Applying ML to the practice of law involves a plethora of legal services, such as e-disclosure, predictive forensics, assessment of evidence, case law analysis, argumentation mining, analysis of applicable law, and quantitative legal prediction. Much of its success is due to a variety of ML techniques that perform natural language processing, seeking to develop statistically accurate relationships between an input (documents that are potentially relevant for evidence, case law, legal briefs or memos, doctrinal text, legislation, and other types of regulation) and a desired output (relevant documents, relevant lines of argument, precedent, doctrine, applicable legislation, or regulation). To the extent that the techniques are used for prediction, the output will mostly concern the outcome of upcoming judgments, decisions of public administration, or even the outcome of judgments by particular judges or administrators.⁵⁴

simulate this – whether such equilibrium is reached and how the algorithms should be trained to accomplish this will necessarily depend on human judgment.

⁵¹ Rosalind Picard, *Affective Computing* (Cambridge, MA: MIT Press, 1997).

⁵² Ran R Hassin, James S Uleman, & John A Bargh, *The New Unconscious*, Oxford Series in Social Cognition and Social Neuroscience (Oxford: Oxford University Press, 2005).

⁵³ Mireille Hildebrandt, 'The Force of Law and the Force of Technology' in MRP McGuire & Thomas J Holt, eds, *The Routledge International Handbook of Technology, Crime and Justice* (New York: Routledge, 2017) 579.

⁵⁴ Katz, 'Quantitative Legal Prediction,' supra note 1; Nikolaos Aletras et al, 'Predicting Judicial Decisions of the European Court of Human Rights: A Natural Language Processing Perspective' (2016) 2 Peer Journal of Computer Science e93; Daniel Martin Katz, Michael J Bommarito II, & Josh Blackman, 'A General Approach for Predicting

The shift from law as information to law as computation raises a number of issues when it comes to legal intelligence. In light of the concerns outlined above, I will discuss four implications that may disrupt the concept and the Rule of Law: (a) the opacity of ML software may render decisions based on its output inscrutable and thereby incontestable; (b) the shift from meaningful information to computation entails a shift from reason to statistics and from argumentation to simulation; (c) a set of fundamental rights may be infringed, compromised, or even violated in the process of developing and testing data-driven legal intelligence, notably the right to privacy, to non-discrimination, to the presumption of innocence, and to due process, while also impacting consumer and employee protection and competition law. Finally, I argue that (d) to the extent that the algorithms become highly proficient, due to being trained by excellent domain experts in law, lawyers may outsource part of their work, as a result of which they may deskill as the software achieves high levels of accuracy. At some point, it may be difficult for these lawyers to check whether the software ‘gets it right,’ confronting us with a new Catch-22. The latter may be the most foundational implication of artificial legal intelligence.

The opacity of ML operates at different angles.⁵⁵ First, the opacity can be due to deliberate hiding of the source code and the relevant training and test data, as they may be part of the trade secret or intellectual property rights of a law firm or the technology firm that developed the software. After having invested in the algorithms, which may have been trained by expensive legal experts, companies may not wish to disclose the inner workings of their systems. Such deliberate opacity raises issues with the contestability of the output of the software, rendering it difficult or even impossible to test its reliability and to investigate the assumptions that inform the development of the algorithms. Software verification is a sub-discipline of computing science, engaging in either mathematical or empirical verification (preferably both).⁵⁶ It can verify mathematical

the Behavior of the Supreme Court of the United States’ (2017) 12:4 PLOS One 12 e0174698.

⁵⁵ Jenna Burrell, ‘How the Machine “Thinks”: Understanding Opacity in Machine Learning Algorithms’ (2016) 3 Big Data and Society 1; Mireille Hildebrandt, ‘Profile Transparency by Design: Re-Enabling Double Contingency’ in Mireille Hildebrandt & E de Vries, eds, *Privacy, Due Process and the Computational Turn* (Abingdon: Routledge, 2013) 221; Bert-Jaap Koops, ‘On Decision Transparency, or How to Enhance Data Protection after the Computational Turn’ in Mireille Hildebrandt & E de Vries, eds, *Privacy, Due Process and the Computational Turn* (Abingdon: Routledge, 2013) 196.

⁵⁶ Arvind Ramanathan et al, ‘Integrating Symbolic and Statistical Methods for Testing Intelligent Systems: Applications to Machine Learning and Computer Vision’ (2016) 2016 Design, Automation Test in Europe Conference Exhibition 786.

proof of the system by inquiring into the hypotheses space, checking for potential over-fitting or over-generalization, and measuring the relevance of the optimization in view of the data used and the given purpose. If no access is given to all of this, one could still do empirical testing of the input and the output, irrespective of how the system came to its results. If the software and/or the training data are trade secrets, however, it will not be easy to perform such empirical testing as too much information may be missing. This warrants a legal obligation to work with open source software and to provide access to specified watchdogs capable of testing and contesting the software as well as the training and test data. We must take into account that much of this software will be developed by big technology firms such as IBM and a myriad of start-ups, not necessarily by practising lawyers, meaning there is no vested interest in, or experience with, the issues of the Rule of Law.⁵⁷

Second, the opacity may be due to the fact that neither domain experts in law nor those subject to law have developed any skills to scrutinize these systems. Whereas most of us have learned to read and write, we were not trained to ‘read’ and ‘write’ statistics, we cannot argue against the assumptions that inform ML applications, and we miss the vocabulary that frames ‘training sets,’ ‘hypotheses space,’ ‘target functions,’ ‘optimization,’ ‘over-fitting,’ and more. Thus, even if the experts manage to verify the software, most of us lack the skills to make sense of it; we may in fact be forced to depend on claims made by those who stand to gain from its adoption and on the reputation of those offering this type of data-driven legal service. This is also the case when we buy and drive a car, but the reliability here is easier to test. We recognize a car crash and would not appreciate a car that drives us from A to C if we want to get to B; with legal intelligence, we may simply not detect incorrect interpretations.

The shift from reason to statistics and from argumentation to simulation is the most obvious issue for those subject to law and should raise awareness with legal practitioners and scholars. Quantified legal prediction (QLP) can only simulate legal decision making, it processes signs (data), not meaning (content).⁵⁸ The inferences it makes are based on mathematical functions that optimize the mathematical relationship between data points in legal text (or data points regarding specific judges, courts, or administrators) and data points in other legal text (for instance, a variable that depicts whether an injunction is provided, a violation is confirmed, or compensation is determined). QLP does not argue, as lawyers

⁵⁷ ‘Legal Tech Market Map: 50 Startups Disrupting the Legal Industry’, *CB Insights* (13 July 2016), online: CBInsights <<http://www.cbinsights.com/research/legal-tech-market-map-company-list/>>.

⁵⁸ Katz, ‘Quantitative Legal Prediction,’ *supra* note 1.

do; it measures and processes elements of their usage of natural language, thus providing feedback on potentially relevant patterns. Though it feeds on legal argumentation and decision making, it does not add its own arguments, merely providing numerical re-articulation and mathematical functions to connect the dots. To ‘read’ the relevant statistics, one would have to break the spell of the second type of opacity, as discussed above, and develop the adversarial skills of ‘speaking’ statistics. This does not necessarily imply less of Michel Foucault and Ronald Dworkin and more of Shannon and Richard Susskind, as Daniel Katz suggests,⁵⁹ but it should suggest, for instance, that the different schools of statistics generate different outcomes that will make a real difference, especially once nearly anything ‘everywhere’ turns into a data-driven Internet of things.⁶⁰

Legal education should make sure that ‘reading’ statistics is an augmentation of human acuity and a contribution to an extended human reason, not its replacement. Neither lawyers nor those subject to law should become statisticians, but both should learn to speak the language of variables, functions, correlations, training sets, hypotheses space, while getting the hang of supervised and unsupervised ML algorithms. This should ultimately help us to speak law to the power of statistics, by which I do not mean to take a silly moralistic approach to ML but, rather, a serious understanding of where the bias inherent in ML applications becomes relevant in terms of the law – for instance, when indirect prohibited discrimination is detected, when disproportional processing of personal data infringes on private life, or when unwarranted monitoring of groups (smart policing, predictive forensics) meddles with the presumption of innocence.⁶¹ This should resolve any tendency to replace legal reasoning with statistical computation, while acknowledging that statistics may contribute to better informed legal reasoning – if done well.

⁵⁹ Daniel Martin Katz, ‘The MIT School of Law? A Perspective on Legal Education in the 21st Century’ (2014) 5 U Ill L Rev 1431 [Katz, ‘MIT School of Law’]. On the lack of methodological integrity in the use of statistics in social science, see e.g. D Mayo, ‘Gigerenzer at the PSA: “How Fisher, Neyman-Pearson, and Bayes Were Transformed into the Null Ritual”: Comments and Queries (ii)’ (blog), *Error Statistics Philosophy* (8 November 2016), online: Error Statistics Philosophy Blog <<https://errorstatistics.com/2016/11/08/gigerenzer-at-the-psa-how-fisher-neyman-pearson-bayes-were-transformed-into-the-null-ritual-comments-and-queries-i/>>, though I am all for serious introductions to the seminal work of Shannon, Wiener, and Gigerenzer. For an interesting example of adversarial statistics, see JASP, online: JASP <<https://jasp-stats.org>>.

⁶⁰ Adam Greenfield, *Everyware: The Dawning Age of Ubiquitous Computing* (Berkeley: New Riders, 2006).

⁶¹ Solon Barocas, Sophie Hood, & Malte Ziewitz, ‘Governing Algorithms: A Provocation Piece,’ SSRN (29 March 2013), online: SSRN <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2245322>; Solon Barocas & Andrew D Selbst, ‘Big Data’s Disparate Impact’ (2016) 104 Cal L Rev 671.

Quasi-invisible infringements of human rights and other types of legal protection will be enabled if legal intelligence is taken for granted, notably when it is meant to replace, instead of inform, legal enactment. We should not fool ourselves, however, by taking for granted that current legal practice entails no such infringements. If legal intelligence is designed and employed to visualize previously invisible discrimination, unsubstantiated conclusions, missing facts, hidden arguments, or arbitrary decision making that flies in the face of both legal certainty and justice as fair and equal treatment, we should applaud it and learn how to 'read' it as such. We should, however, not be naive about the design of legal intelligence. Other than Larry Ribstein, as quoted by Katz, I am not sure that the 'harsh winds of the market' will necessarily blow in the right direction, considering who will be funding the architecture of legal intelligence.⁶² If IBM's robolawyer (as their application of legal intelligence has been framed), is trained by top lawyers from major corporate law firms, it will reduplicate their concerns and those of their clients.⁶³ If the contracts between IBM and such law firms give IBM a hold on the data and/or require law firms to recognize copyright or patent on the software on the side of IBM, we should expect to be confronted by some of the drawbacks that plague other oligarchies driven by large technology companies.

2 *the new catch-22: judgmental atrophy*

Both the purported opacity of some types of ML and the schism between mathematical relations and the semantics of natural language imply that the outsourcing of legal intelligence to machines raises the problem of how to evaluate if and how these machines are actually 'getting it right.' In his chilling *The Glass Cage*, Nicholas Carr develops a persuasive argument that while smart machines get better in simulating the skills of their human trainers, domain experts may fall prey to gradual deskilling.⁶⁴ In a sense, this follows directly from John Searle's famous 'Chinese room argument,' which demonstrates that while Alan Turing's so-called 'thinking machines' may function as semi-perfect simulators, they never get down to the 'real thing.'⁶⁵ They lack consciousness and self-consciousness, and even if these were to be epiphenomena of biological evolution,⁶⁶ they

⁶² Katz, 'MIT School of Law,' supra note 59 at 101.

⁶³ Jason Koebler, 'Rise of the Robolawyers,' *The Atlantic* (April 2017), online: The Atlantic Monthly Group <<https://www.theatlantic.com/magazine/archive/2017/04/rise-of-the-robolawyers/517794/>>.

⁶⁴ Nicholas Carr, *The Glass Cage: Automation and Us* (New York: WW Norton, 2014).

⁶⁵ John Searle, 'Minds, Brains, and Programs' (1980) 3:3 *Behavioral and Brain Sciences* 517.

⁶⁶ Daniel C Dennett, *The Intentional Stance*, rev ed (Cambridge, MA: MIT Press, 1989) at 300.

make ‘a difference that makes a difference’ compared to us and to a law that is based on generating both meaning and legal effect.⁶⁷

In a remarkable article, Federico Cabitza develops a similar argument in relation to medical expertise that is trained into diagnostic ML applications.⁶⁸ He narrates how an ML system (Zeb) that is capable of ‘reading’ radiological images achieves a higher performance than the human radiologists that trained the algorithms, getting it right 85–98 per cent of the time. Cabitza raises the question what happens if radiologists routinely consult Zeb, finding that they agree with its diagnosis most of the time. He suggests:

Maybe Zeb would help novices learn how to interpret both easy and difficult images; improve the residents’ skills more quickly also by challenging them with tricky simulations based on real cases; and teach also expert radiologists how to solve cases that before its arrival would be too difficult for a fast analysis and require several meetings, or worse yet, further examinations for the patient. Conversely, Zeb would perhaps undermine the self-confidence of the expert radiologists, . . . after a few times that their interpretation[s] differ from Zeb’s and . . . proved to be the correct one. Likewise, it could make the young or less brilliant radiologists more lazy and dependent on its recommendations. Thus the main point here is whether systems like Zeb have the potential to actually deskill or ‘spoil’ the physicians in the long run.

It seems important to stress that Zeb could actually improve the accuracy of medical diagnosis, while acknowledging that in the long run it could deskill the human specialist.

Cabitza enumerates the following unintended consequences in relation to ML-driven decision support systems (MLDSS). He speaks of ‘empirical anopsia and conversational hypoacusia, which both regard the representation of the available Experience (E) and the input of the MLDSS.’ He describes ‘probabilistic tinnitus, which regards the representation of the outcome of the task (T) and hence the output of the MLDSS’ and what he calls ‘epistemic sclerosis, which regards the representation of the

⁶⁷ G Bateson, *Steps to an Ecology of Mind* (New York: Ballantine, 1972) at 459. This is where Solum seems to get carried away, suggesting that machines will produce artificial meaning, just like groups of people do. He fails to register that even if machines could be said ‘to produce’ meaning, it is human beings that attribute such meaning to the outcome of machine operations. I would suggest instead that human beings ‘naturally’ produce artificial meaning, while machines ‘artificially’ produce a kind of ‘naturalized’ meaning (emulating what they infer as the state of the art in legal reasoning). Such ‘naturalization’ comes close to Alarie’s legal singularity, again mistaking a simulation for what is simulated. See Solum, ‘Artificial Meaning,’ *supra* note 1.

⁶⁸ Federico Cabitza, ‘Breeding Electric Zebras in the Fields of Medicine,’ Doc ArXiv:1701.04077 [Cs] (15 January 2017), online: Cornell University Library <<http://arxiv.org/abs/1701.04077>>.

performance of the task itself.’ Finally, he refers to ‘metric paroxysm, that regards the very way performance is assessed (that is, P)’ and ‘semiotic desensitization, judgmental atrophy and oracular rush which derive from relying too much on the “accurate” output of the MLDSS and its effect on the interpretative work of medical doctors,’ and, lastly, he discusses ‘metric paroxysm, that is an ill-grounded trust in numeric measures of the performance of ML-based classifiers’ – for instance, when ‘the value of a ML model emerges more from a complex trade-off between its accuracy and its explainability, which must be evaluated qualitatively, rather than by means of a single figure.’⁶⁹ Cabitza combines a hilarious depiction of the lack of proper judgment in terms reminiscent of learned medical diagnosis with an inspiring acuity about the drawbacks of human reliance on the automation of high-level cognitive tasks, adding that an

ML approach risk[s] to freeze into the decision model two serious and often neglected biases: selection bias, occurring when training data (the above experience E) are not fully representative of the natural case variety due to sampling and sample size; and classification bias, occurring when the single categories associated by the raters to the training data oversimplify borderline cases (i.e., cases for which the observers do not agree, or could not reach an agreement), or when the raters misclassify the cases (for any reason). In both cases, the decision model would surreptitiously embed biases, mistakes and discriminations and, worse yet, even reiterate and reinforce them on the new cases processed.⁷⁰

What is at stake here is the ability of doctors to recognize when the automation goes awry, misrepresents relevant cases, or misinterprets relevant causation. The risk of automating incorrect, imprecise, or debatable judgment is not restricted to medical diagnosis. Similar caveats apply when training the algorithms of artificial legal intelligence, where interpretation and subsequent qualification define and determine the legal effect of human intercourse. We need to train our machines to augment human legal intelligence, not to replace it: ‘[U]ltimately, this technology is only as good as the training it receives. A poorly trained API [application programming interface] will consistently produce erroneous results, while a well-trained API will reliably generate accurate results.’⁷¹ As law evolves, the software should evolve, and this is where lawyers must come in to make sure fundamental legal principles are taken into account; otherwise ‘self-driving law’ may decide on shortcuts that we appreciate in

⁶⁹ Ibid.

⁷⁰ Ibid.

⁷¹ Yoon, ‘Post-Modern Lawyer,’ *supra* note 1 at 467, n 48. An application-programming interface, for instance, allows one to build new applications on existing platforms.

self-driving cars but not in the decision-making infrastructure of human society.

IV *Concluding remarks: from 'legal by design' to 'legal protection by design'*

In this article, I have made an attempt to critique some of the rosy scenarios of artificial legal intelligence, notably where they seek to emulate a 'complete law' that 'closes the gaps in the law' to achieve 'legal singularity,' finally making 'legal uncertainty obsolete.'⁷² Though much can be learned from such scenarios, they tell the tales of a totalitarian system, nourishing on a supposedly complete prediction of human behaviour, enabling an approximation of complete compliance, envisaging human behaviour that is 'legal by design.'⁷³ My argument boils down to the fact that this is not about law but, rather, about discipline or a nefarious type of public administration. Does this mean that we should forego artificial legal intelligence? As Alarie rightly foresees, current market incentives will probably drive artificial intelligence to a point where the law will be more complex than we can even imagine.⁷⁴ Nevertheless, I believe that, as lawyers, we should engage with artificial legal intelligence to make sure it aligns with law and the Rule of Law in a testable and contestable way. As a matter of fact, this may save us from ridiculously complex systems outcompeting each other with regard to big business compliance issues.

How to bring artificial legal intelligence under the Rule of Law is not obvious. It will require a specific design of the upcoming computational architecture of our legal systems; indeed, it requires us to reinvent law and the Rule of Law, setting the right defaults, developing the right standards, translating fundamental legal principles into the hardware, the operating systems, the firmware, the software, the applications, and the ML methodologies we are on the verge of embracing. In my work on privacy, non-discrimination, and due process, I have coined this 'legal protection by design.' Whereas 'legal by design' is all about compliance, 'legal protection by design' aims to safeguard our ability – as individuals – to challenge automated decision systems, by providing time and space to test and contest the workings of such systems. This will not be an easy task for the reasons given earlier in this article. We may also ponder the wisdom of catalyzing the legal arms race between big business and big

⁷² Terms taken from Alarie, 'Path of the Law,' *supra* note 1 at 452, 451, 445.

⁷³ Paul Lippe, Daniel Martin Katz, & Dan Jackson, 'Legal by Design: A New Paradigm for Handling Complexity in Banking Regulation and Elsewhere in Law' (2015) 93:4 Oregon L Rev 831.

⁷⁴ Alarie, 'Path of the Law,' *supra* note 1 at 451.

regulators, but I believe that there are ways of getting this right, emulating our own legal intelligence in conjunction with ML systems to clarify hidden assumptions, missed arguments, debatable facts, and prohibited biases. That way we can speak law to the power of statistics, based on that same power. For those familiar with the paradox of the Rule of Law that is a familiar challenge.