

Research Study on Computational Law

Pauline McBride and Laurence Diver

[DRAFT]

Cite as:

Pauline McBride and Laurence Diver, *Research Study on Computational Law* (Brussels 2023), funded by the ERC Advanced Grant 'Counting as a Human Being in the Era of Computational Law' (COHUBICOL) by the European Research Council (ERC) under the HORIZON2020 Excellence of Science program ERC-2017-ADG No 788734 (2019-2024)

Table of contents

1	Introduction: foundations and scope	1
1.1	<i>Introduction</i>	1
1.2	<i>Computational transformations</i>	1
1.2.1	A market-oriented perspective	1
1.2.2	A deeper perspective – the contingencies of law-as-we-know-it	2
1.3	<i>Law’s mode of existence</i>	3
1.4	<i>Anticipating the impact of computational legal technologies</i>	4
1.5	<i>Conceptual tools and approaches</i>	6
1.5.1	Gibson’s concept of ‘affordance’	6
1.5.2	Beyond affordances: complementary perspectives	8
2	The impact of data-driven technologies	10
2.1	<i>The rise of data-driven legal technologies</i>	10
2.2	<i>The affordances of data-driven technologies</i>	13
2.2.1	Answering an objection	15
2.3	<i>Data-driven technologies as change agents and influencers</i>	18
2.3.1	Making way for a new normativity	18
2.3.2	Engines of influence	20
2.3.3	New seats of power	23
2.3.4	Summary	24
2.4	<i>A closer inspection of agentic effects: engines of influence, seats of power</i>	25
2.4.1	Engines of influence, seats of power	25
2.5	<i>A closer inspection of agentic effects: making way for data-driven normativity</i>	27
2.5.1	The texture of data-driven normativity	28
2.5.2	The implications of data-driven normativity: legal protection and the Rule of Law	29
2.6	<i>Conclusion</i>	34
3	The impact of code-driven legal technologies	36
3.1	<i>Introduction</i>	36
3.2	<i>Rules as Code</i>	37
3.2.1	The COHUBICOL lens	39
3.2.1.1	Legality and legalism	39
3.3	<i>The texture of code-driven normativity</i>	41
3.3.1	Mixing legal and technological normativity	43
3.4	<i>A spectrum of impact on law</i>	45

3.4.1	Stage 1: Legal metadata and machine-readable documents	45
3.4.1.1	Structural markup	46
3.4.1.2	Drafting software	47
3.4.1.3	Potential impact(s) of code-driven law	48
3.4.2	Stage 2: policy development and automated analysis of rules	50
3.4.2.1	Logic checking	51
3.4.2.2	Policy development and parallel drafting	52
3.4.2.3	Potential impact(s) of code-driven law	54
3.4.3	Stage 3: bespoke languages and digital-first laws	55
3.4.3.1	RegelSpraak: a controlled natural language	56
3.4.3.2	Catala: a domain-specific language	57
3.4.3.3	A 'single source of truth'	58
3.4.3.4	Digital-first laws	59
3.4.3.5	Potential impact(s) of code-driven law	60
3.5	<i>Anticipating legal protection under code-driven law</i>	61
3.5.1	Who benefits?	61
3.5.2	Increased complexity and maintenance	62
3.5.3	Interpretative authority	63
3.5.3.1	The mirage of human-readable code	64
3.5.3.2	Technological normativity and interpretation	65
3.5.3.3	Technological normativity and discretion	66
3.5.3.4	Formalisation and the shaping of policy	68
3.5.4	RaC and the effect on legal effect	70
3.5.4.1	The mode of existence of code-driven law	70
3.6	<i>Conclusion: treading the line</i>	73

1 Introduction: foundations and scope

Pauline McBride

1.1 Introduction

In the COHUBICOL Project proposal we read that

the core thesis of the research is that the upcoming integration of computational law into mainstream legal practice, could transform the mode of existence of law and notably of the Rule of Law.

In this Research Study we explore the dynamics, modalities and effects of such transformation. Taking the Research-Study on Text-Driven Law¹ as a foundation, we demonstrate how data-driven and code-driven legal technologies have implications for law-as-we-know-it, the Rule of Law and the nature of the protection afforded by law.

1.2 Computational transformations

Increasingly legal technologies are used to carry out tasks historically reserved to lawyers. Sceptics and proponents alike describe these technologies as transformative, disruptive, productive of change.² The sector is booming; it is essential to grasp the implications of such change. What kind of change might we anticipate? What are the drivers of such change?

1.2.1 A market-oriented perspective

For Susskind legal technologies facilitate an evolution of legal services,³ and it is the market which is the locus of transformation.⁴ The market drives a progression from bespoke legal services to commoditisation

¹ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt, 'Research Study on Text-Driven Law (Brussels 2023), Funded by the ERC Advanced Grant "Counting as a Human Being in the Era of Computational Law" (COHUBICOL) by the European Research Council (ERC) under the HORIZON2020 Excellence of Science Program ERC-2017-ADG No 788734 (2019-2024)' (COHUBICOL, 20 September 2023) <<https://www.cohubicol.com/news/publication-of-the-cohubicol-research-study-on-text-driven-law/>> accessed 15 October 2023.

² John O McGinnis and Russell G Pearce, 'The Great Disruption: How Machine Intelligence Will Transform the Role of Lawyers in the Delivery of Legal Services' (2019) 13 *Actual Problems of Economics and Law* <<https://www.rusjel.ru/jour/article/view/25>> accessed 16 August 2023; Susanne Chishti (ed), *The Legaltech Book: The Legal Technology Handbook for Investors, Entrepreneurs and FinTech Visionaries* (John Wiley & Sons 2020); Mireille Hildebrandt, *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology* (Paperback edition, Edward Elgar Publishing 2016).

³ Richard Susskind, *The End of Lawyers?: Rethinking the Nature of Legal Services* (Revised edition, Oxford University Press 2010) 28–33. Mark A Cohen, "'Legal Services' Are Whatever Buyers Need To Solve Business Challenges' (*Forbes*) <<https://www.forbes.com/sites/markcohen1/2019/03/03/legal-services-are-whatever-buyers-need-to-solve-business-challenges/>> accessed 8 November 2023. The language of 'legal services' already signals a particular perspective on what law is and what it is for.

⁴ Susskind (n 3) chs 2-4.

of those services.⁵ This progression, Susskind argues, is good for the consumer of legal services, whether lawyers like it or not.

Susskind and Susskind develop this theme in later work. The endpoint of commoditisation is presented as ‘externalisation’ of services.⁶ Machine learning systems will ‘generate practical expertise’ independently from the humans who design or use them.⁷ This vision of the future of law is one in which the remit of lawyers is diminished and the role for machines is increased. Tasks which were the preserve of lawyers are handled by machines.⁸

The language of commoditisation is apt to conceal some of the implications of use of legal technologies to carry out tasks that would previously have been carried out by lawyers. It tends to suggest that while the mode of delivery of these tasks is different, the output might be (more or less) the same; it suggests, that is, that commoditisation (and externalisation) is the transformation. It downplays the agentive role of technologies, the institutional dimension of law as a practice, the wider systemic effects of such change.

Like Susskind we believe that legal practice will be transformed, and that data driven technologies will have a part to play in that transformation. However, we consider that the levers and dynamics of transformation are both more subtle and more far-reaching in their effects than Susskind suggests. There is more at stake than the scope of tasks that are reserved to lawyers. Even on Susskind’s account the handover of tasks to machines entails the establishment of new practices and the creation of new seats of power.

1.2.2 A deeper perspective – the contingencies of law-as-we-know-it

The Project’s Research Study on Text-driven law re-articulates Hildebrandt’s conception of law’s ‘mode of existence’ of law – a concept introduced by Souriau and developed by Latour.⁹ The reference to a ‘mode of existence’ signals that the way in which law exists now, along with the protection it affords, is not a given. For Hildebrandt, law has not always existed in its current mode, nor need it do so in the future. However, the concept of law’s ‘mode of existence’ does more than highlight the contingency of law. It is a frame for exploring and exposing the nature of those contingencies.

⁵ *ibid* 28–33. McGinness and Pearce similarly frame the disruptive effects of legal technologies in terms of ‘commoditisation’ and impacts on the legal ‘market’. McGinnis and Pearce (n 2) 3042, 3054. Munisami suggests that there is agreement in the literature that ‘automation will shift major parts of the legal work lawyers typically do towards “commoditization”’. Kayal Munisami, ‘Legal Technology and the Future of Women in Law’ (2019) 36 *Windsor Yearbook of Access to Justice* 164, 166.

⁶ Richard E Susskind and Daniel Susskind, *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (First edition, Oxford University Press 2015) 196, 197.

⁷ *ibid* 107, 267.

⁸ Hildebrandt notes that ‘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.’ Mireille Hildebrandt, ‘Law as Computation in the Era of Artificial Legal Intelligence: Speaking Law to the Power of Statistics’ (2018) 68 *The University of Toronto Law Journal* 12, 28.

⁹ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 1) 13.

1.3 Law's mode of existence

As a preliminary it is useful to tease out what Latour grapples with and hopes to convey by 'modes of existence'. Latour wishes to debunk the 'modernist' idea that the world exists in and can be grasped from a single perspective. Instead, he presents a view of the world 'as a congeries of perspectives or points of view grounded in different modes of existence'.¹⁰ Latour maintains that the different institutions of science, law and religion, for example, have different ways of producing truth.¹¹ Indeed Latour, in a move which mediates between being and knowing, ontology and epistemology, equates a mode of existence with a truth regime or a 'regime of veridiction'.¹²

Significantly, for Latour, law, like religion and politics, is also quintessentially a mode of enunciation.¹³ These various framings of law as mode of existence, regime of veridiction and mode of enunciation already hint at the reflexiveness of law's normativity. That is, law's mode of existence, what makes it 'law', is both the 'legal' character of its enunciations and the means by which those enunciations are verified as 'legal'.

For Latour, Austin's concept of speech acts is the key to understanding how law 'is', 'utters' and 'verifies' in a mode which is 'legal'.¹⁴ In Austin's terminology, enunciations are speech acts which 'do things'; they have performative effect. A speech act will have this effect if and when what Austin terms their 'conditions of felicity' are satisfied. For example, I name a ship if and only if my act of naming meets the conditions for that utterance to take effect, so qualifying as a speech act.¹⁵ What sets law apart from other modes of enunciation is that the speech acts of law are performative when they produce binding effects for legal subjects. This is true when the speech acts meet the 'felicity conditions' articulated by law for those utterances to clothe some set of facts, action or norm with specified legal effects.¹⁶ This reflects the tautological character of law: law both is and specifies its own regime of veridiction.¹⁷ For this reason, we can view law as both a collection of speech acts with the capacity to bind, and as involving the *projection* of law to encompass new facts and actions and produce new effects by the operation of speech acts. In Latour's words, these aspects of law 'are the same ... but simply taken at different moments in their process of crystallization'.¹⁸ The first views law statically (law simply 'being'), the second

¹⁰ Laurent de Sutter, 'Plasma! Notes on Bruno Latour's Metaphysics of Law' in Kyle McGee (ed), *Latour and the Passage of Law* (Edinburgh University Press 2015) 202.

¹¹ Bruno Latour, *The Making of Law: An Ethnography of the Conseil d'Etat* (Polity 2010) 8.

¹² Bruno Latour, 'Biography of an Inquiry: On a Book about Modes of Existence' (2013) 43 *Social Studies of Science* 287, 287.

¹³ Bruno Latour, *An Inquiry into Modes of Existence: An Anthropology of the Moderns* (Harvard University Press 2013) 375.

¹⁴ *ibid* 56.

¹⁵ JL Austin, *How To Do Things With Words* (2nd ed, Harvard Univ Press, 1975) 5, 23, 37.

¹⁶ In Latour's words 'With law, characters become assigned to their acts and to their goods. They find themselves responsible, guilty, owners, authors, insured, protected.' Latour, *An Inquiry into Modes of Existence* (n 13) 370.

¹⁷ As Pottage explains, 'law comes into being paradoxically, as an effect of the identification of certain enunciations or transactions as "legal" by reference to a criterion that is posited by those enunciations or transactions themselves.' Alain Pottage, 'The Materiality of What?' (2012) 39 *Journal of Law and Society* 167, 173.

¹⁸ Bruno Latour, 'The Strange Entanglement of Jurimorphs' in Kyle McGee (ed), *Latour and the Passage of Law* (Edinburgh University Press 2015) 343.

draws attention to law in action, or ‘throw[ing] itself forward’ through the authoritative ascription of legal effect.¹⁹

Law’s very existence, therefore, is bound up with speech acts, with words that do things, with a very specific kind of performativity, with ‘felicity conditions’ that are themselves set by law. Such is the nature of law that one cannot even ‘speak of the law without speaking *legally*.’²⁰ Speaking of law, speaking ‘legally’, involves speaking ‘within’ the law. Latour, in short, draws attention to the fact that law’s mode of existence depends on ways of speaking – and reasoning – that are distinctively ‘legal’.²¹

1.4 Anticipating the impact of computational legal technologies

Latour’s focus on law’s distinctive regime of enunciation is such that the legal anthropologist Alain Pottage complains:

Latour’s analysis of law often proceeds as though there were actually nothing more to law than a process of enunciation. The effect is to suggest that law is not a material world in the same sense as science or technology.²²

Materiality is present in Latour’s account of the mode of existence of law, but it has a supporting, not a leading role.²³ Hildebrandt, by contrast, develops an account of law’s mode of existence which accords greater significance to the materiality of law. Thus, in the Research Study on Text Driven Law we read that:

¹⁹ Bruna Latour, Owen Martell (tr), ‘AIME Platform Vocabulary: “Enunciation, Regime of Enunciation”’ <<http://modesofexistence.org/>>; Latour notes ‘[i]n this inquiry, being and enunciation can almost be taken as synonymous since it is the nature of a being to utter itself, to exist, to transit, to throw itself forward through the HIATUS of existence or expression’. ‘An Inquiry Into the Modes of Existence’ <<http://www.modesofexistence.org>> accessed 15 October 2023. See also Serge Gutwirth, ‘Providing the Missing Link: Law after Latour’s Passage’ in Kyle McGee (ed), *Latour and the Passage of Law* (Edinburgh University Press 2015). Gutwirth distinguishes between ‘law ... as an intertwined whole of statutes, rules and regulations ...[and] as decision-making or as a practice that produces solutions.’

²⁰ Bruno Latour, ‘The Strange Entanglement of Jurimorphs’ in McGee (n 18) 334.

²¹ See also Serge Gutwirth, ‘Providing the Missing Link: Law after Latour’s Passage’, in McGee (n 18) (‘... the distinctiveness of law lies in the singular mode in which it seizes cases. In other words: everyone can practise law, everyone [who is called to do so] can become a legal practitioner, and that is, when she is moving or moved forward by the legal regime of enunciation ...’); As Latour points out with reference to an anecdote of a child objecting to the theft of their marbles in the language of rights, speaking ‘legally’ is not confined to lawyers. Bruno Latour, ‘Note Brève Sur l’écologie Du Droit Saisie Comme Énociation’ [2004] *Pratiques Cosmopolitiques Du Droit*.

²² Pottage (n 17) 170.

²³ Latour is well aware that law does not circulate in the ether but needs a ‘setup’, a material environment, comprising ‘heterogenous set of elements’ which allows law to circulate. Latour, *An Inquiry into Modes of Existence* (n 13) 32. However, Latour – who popularised the idea that objects can have agency – appears to consider that none of these elements is indispensable. His description of the work of the Conseil d’Etat takes note of the files, books, tables and chairs, the computer database, telephones, paperclips and staplers. However, ‘the nature of Council’, he says, ‘does not depend on its equipment’. Latour, *The Making of Law* (n 11) 167, 168.

The materiality of law can be found in the embodiment of natural language and in the embedding of written law in the technologies of text.²⁴

Of course, speech acts – including the written speech acts of law – depend on the embodiment of natural language.²⁵ Latour also acknowledges that writing has made it easier for law to make linkages between person and speech acts and facts.²⁶ However, unlike Latour, Hildebrandt maintains that the material environment of the ‘information infrastructure’ of law is crucial to law’s mode of existence; law-as-we-know-it is an affordance of text and the printing press.²⁷ Moreover, Hildebrandt’s concern is not for law as a form of discourse *as such* but as a form of discourse that respects human agency and affords legal protection. She is equally concerned about *law as an affordance* of text and the printing press, and legal protection *as an affordance of law*.²⁸

In order to highlight law’s dependence on text and the printing press Hildebrandt traces the implications of the shift from orality to script, from handwritten to printed script both generally and in the context of the legal tradition. The materialisation of legal norms in text creates a separation across time and space between an author of a legal norm (a judge or a legislator) and its audience. It makes it possible for law to be directed to many and across a wide geographic area.²⁹ Such distantiation, in turn, creates a need for interpretation of legal norms, gives rise to the emergence of a legal cadre, allows for the creation of large jurisdictions and the concept of equality under the law.³⁰ The need for interpretation is key; this in turn gives rise to a need for deliberation and ultimately closure of interpretative conflicts. The growing complexity of law eventually creates the conditions necessary for the emergence of the Rule of Law.³¹ Thus:

The specific nature of the technology of the text thus leads a shift from ‘rule by law’, i.e. the law as an instrument by which governments enforce their own interpretation of the norms they issue, to ‘Rule of Law’, i.e. the law as a system of checks and balances that institutes countervailing powers, such that public administration and even the legislature itself are brought under the Rule of Law. In that sense the core principles of the Rule of Law (such as contestability and

²⁴ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 1) 35.

²⁵ Moreover, as van den Hoven notes, ‘... it is important to note speech acts will not magically ‘do what they say’, and legal speech acts will therefore also not automatically bring about the legal protection we desire or need. Crucially, whether a speech act has performative effect therefore depends on a shared acceptance of or acquiescence in the world of institutional facts it is embedded in. It builds on a pragmatist understanding of language and depends ‘on a shared background consisting of hidden assumptions, mutual beliefs and a joint practice that grounds the use and thereby the meaning of words and more generally of human action’. Emilie van den Hoven, ‘Legal Effect, Sources of Law, and Jurisdiction’ in *ibid* 70.

²⁶ Latour, *An Inquiry into Modes of Existence* (n 13) 371.

²⁷ Latour asserts that since Roman times, ‘no radical innovation has altered the art of enunciating the law’. *ibid* 217.

²⁸ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 1) 18. We discuss the concept of an affordance in section 1.5.1.

²⁹ Mireille Hildebrandt, ‘A Vision of Ambient Law’ in Roger Brownsword and Karen Yeung (eds), *Regulating Technologies* (Hart 2008) 184.

³⁰ Mireille Hildebrandt, ‘A Vision of Ambient Law’ in Roger Brownsword and Karen Yeung (eds), *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes* (Hart 2008) 184. See also David Harvey, ‘Law and the Regulation of Communications Technologies: The Printing Press and the Law 1475-1641’ [2005] *Australian and New Zealand Law & History Society E Journal* 160.

³¹ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 1) 4.

accountability) are not merely historical artifacts but also technological artifacts, directly linked to the flexibility of natural language and the responsive autonomy of text-driven normativity.³²

For Hildebrandt, law's mode of existence then is not only contingent on the constant tracing and retracing of connections between speech acts, persons, facts, actions, texts and legal effects³³ – an activity which presupposes and requires a process and practice of legal reasoning – it is also contingent on its information infrastructure. A fundamental change in the material and institutional environment of law has the potential to transform law's mode of existence, bring about the loss or reconfiguration of those affordances that are the hallmark of law-as-we-know-it: contestability, interpretation, closure³⁴ and the Rule of Law – and with them the protection afforded to legal persons by law. It has the potential to alter the mode of existence of law. Much more is at stake, therefore, in the turn to computational law than the loss of lawyers' monopoly over certain tasks. Hildebrandt is astute therefore to urge a 'perspicacious sensitivity to the affordances of the AI that we are constructing ...'³⁵

1.5 Conceptual tools and approaches

It may be supposed that with the benefit of history, the significance of text and the printing press can readily be appreciated and that the transformative effects of current technologies may be rather less easy to discern. There is surely some truth in this statement, but it is worth noting that Hildebrandt's account of the significance of text and the printing press for law-as-we-know-it is neither 'a merely historiographical undertaking [nor] a matter of social scientific research into the causation of modern law and the modern state.'³⁶ Hildebrandt teases out the implications of text and the printing press by reference to their *affordances*. This concept, and the underlying assumptions which inform it, are key to understanding the dynamics of the transformation brought about by and through technologies.

1.5.1 Gibson's concept of 'affordance'

The term 'affordance' was coined by the ecological psychologist James Gibson. Gibson's concept of affordance captures the idea of the action possibilities offered by an environment (including places, objects, people, animals) to an actor (human or not):

The affordances of the environment are what it offers the animal, what it provides or furnishes, either for good or ill. The verb to afford is found in the dictionary, but the noun affordance is not. I have made it up. I mean by it something that refers to both the environment and the animal

³² *ibid.*

³³ See Latour, *The Making of Law* (n 11) 208, 218; Latour, *An Inquiry into Modes of Existence* (n 13) 370, 371; Sutter (n 10).

³⁴ Hildebrandt notes that law imposes 'closure' in a particular way, that is 'unilaterally after having taken the time to explore uncertainties and ambiguities.' Hildebrandt, *Smart Technologies and the End(s) of Law* (n 2) 183. See also Mireille Hildebrandt, *Law for Computer Scientists and Other Folk* (Oxford University Press 2020) ch 11 <<https://global.oup.com/academic/product/law-for-computer-scientists-and-other-folk-9780198860884?cc=be&lang=en>> accessed 10 August 2019.

³⁵ Hildebrandt, *Smart Technologies and the End(s) of Law* (n 2) 167.

³⁶ *ibid* 177.

in a way that no existing term does. It implies the complementarity of the animal and the environment.³⁷

For example, a tree affords shade and shelter, a chair affords sitting, a stair affords climbing. These action possibilities are not *properties* of the environment or actor(s) (human or otherwise) but arise from the relations between the two. Thus ‘while affordances belong to neither subject nor object, they are potentialities that exist in the world and can do something in it, implying that objects have a certain kind of agency or effectivity.’³⁸ When objects are introduced into an environment, they change the action possibilities for actors. Moreover, since actors and environments are fundamentally (as it were, *symbiotically*) linked,³⁹ such change may in turn be productive of change. To paraphrase and extend Maier and Fadel,⁴⁰ we might say that:

$$\Delta \text{ environment} \Rightarrow \Delta \text{ affordance} \Rightarrow \Delta \text{ behavior}$$

A change in the environment brings about changes in the affordances of the environment, which may in turn bring about changes in the behaviour of actors in that environment, though the interactions, in reality, are not as mono-directional as the representation might imply.⁴¹ The core insight however is that the introduction of technologies into an environment implies the possibility of change.

Gibson’s concept of affordance has considerable explanatory power.⁴² However, his analysis is largely confined to relatively straightforward animal/environment interactions.⁴³ Later scholarship has extended both the concept and the analysis to cater for more complex interactions.⁴⁴ Some of these conceptions diverge from Gibson’s original concept in important respects. Nevertheless, the focus on the relational aspect (animal/environment; object/actor), the idea that both are entangled in way that effects change, remains.

Hildebrandt extends Gibson’s concept of affordance in two respects. First, the concept is extended to capture the way technology can effect change in our institutional environment, specifically how text and

³⁷ James J Gibson, *The Ecological Approach to Visual Perception: Classic Edition* (Psychology Press 2015) 119.

³⁸ Julka Almqvist and Julia Lupton, ‘Affording Meaning: Design-Oriented Research from the Humanities and Social Sciences’ (2010) 26 *Design Issues* 3, 13.

³⁹ Thus, Gibson insists that ‘animal and environment make an inseparable pair. Each term implies the other. No animal could exist without an environment surrounding it. Equally, although not so obvious, an environment implies an animal (or at least an organism) to be surrounded.’ Gibson (n 37) 4.

⁴⁰ Jonathan Maier and Georges Fadel, ‘Affordance Based Design: A Relational Theory for Design’ (2009) 20 *Res. Eng. Des.* 13 The authors refer to ‘structure’ rather than the environment.

⁴¹ See also Roope Oskari Kaaronen, ‘Steps to a Sustainable Mind: Explorations into the Ecology of Mind and Behaviour’ (Helsingin yliopisto 2020) 58 <<http://hdl.handle.net/10138/319046>> accessed 12 November 2023 (suggesting that behavioural change alters the ‘landscape of affordances’).

⁴² However, affordance theory is not without its detractors. Burlamaqui and Dong claim that Gibson’s ‘broad explanation ... is not sufficient for its application by the design community.’ Leonardo Burlamaqui and Andy Dong, ‘The Use and Misuse of the Concept of Affordance’ in John S Gero and Sean Hanna (eds), *Design Computing and Cognition ’14* (Springer International Publishing 2015) 297. Davis and Chouinard refer to various shortcomings. Jenny L Davis and James B Chouinard, ‘Theorizing Affordances: From Request to Refuse’ (2016) 36 *Bulletin of Science, Technology & Society* 241.

⁴³ Oliver says of Gibson that ‘technology, if mentioned, is stone-age’. Martin Oliver, ‘The Problem with Affordance’ (2005) 2 *E-Learning and Digital Media* 402, 404.

⁴⁴ Oliver maintains that these developments devalue the concept. Oliver (n 43).

the printing press afford law-as-we-know-it (law as an affordance of a particular information infrastructure). Second, she extends the concept to capture the affordances of the material and information infrastructure that is law-as-we-know-it (*the affordances of law*).⁴⁵ These conceptual innovations combine Gibson's attention to the material environment with an appreciation that the *institutional* environment of human beings is both material *and* materialised through the embodiment of language. They also involve a recognition that the material and institutional environment of law, in Gibson's language, provides human beings with a form of 'niche' which offers a particular way of life.⁴⁶

The concept of 'affordance' employed by Hildebrandt stays true to Gibson insofar as it 'attend[s] to the ways people and things co-constitute each other'.⁴⁷ It provides a conceptual tool that allows us to anticipate the transformative effects of legal technologies on the material and institutional environment of law, account for change and reflect on how technologies might be designed to conduce to certain values or ends.⁴⁸ Nevertheless, other complementary perspectives also shed light on the transformative potential of technologies.

1.5.2 Beyond affordances: complementary perspectives

In the 'modernist' perspective critiqued by Latour, a sharp division can be made between subjects and objects, persons and things. Persons have agency, things do not. Subjects put objects to use; they alone determine how objects will be used. Objects are 'neutral', without influence. At the other end of the spectrum technological determinism holds that technologies drive and dictate societal change. Gibson's concept of affordances, on the other hand, speaks to the interdependence of persons and things and the possibility for things, as well as persons, to have effects in the world.

Gibson is far from alone in recognising the productive dynamics of relations between persons and things. Latour's actor-network theory emphasises how objects shape action.⁴⁹ Ihde's post-phenomenological account may be understood primarily as an exposition of how technologies shape our perceptions, our

⁴⁵ Hildebrandt, *Smart Technologies and the End(s) of Law* (n 2); M Hildebrandt, 'Law as an Affordance: The Devil Is in the Vanishing Point(s)' [2017] <http://cal.library.utoronto.ca/index.php/cal/article/view/28154/20742> <<https://repository.uibn.ru.nl/handle/2066/173203>> accessed 13 February 2019; Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 1) 18.

⁴⁶ Regarding Gibson's use of the concept of 'niche' in relation to affordances see Gibson (n 37) 120, 121, 133, 135; Erik Rietveld and Julian Kiverstein, 'A Rich Landscape of Affordances' (2014) 26 *Ecological Psychology* 325. See also Laurence Diver, 'Legal Subject, Subjective Rights, Legal Powers' in Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 1) 78 (suggesting that in 'the legal-institutional dimension of the legal ecology, the 'animal' is the legal subject, whose niche consists of the contingent set of rights and powers that it holds.').

⁴⁷ Jenny L Davis, "'Affordances" for Machine Learning', 2023 *ACM Conference on Fairness, Accountability, and Transparency* (ACM 2023) 325 <<https://dl.acm.org/doi/10.1145/3593013.3594000>> accessed 19 October 2023. Hildebrandt explicitly adopts a 'relational' account of law which 'denies that law is independent from its societal, scientific and professional environment, because its existence depends on the performative nature of the social fabric it constitutes and by which it is constituted.' Hildebrandt, *Smart Technologies and the End(s) of Law* (n 2) 172.

⁴⁸ Laurence Diver, *Digisprudence: Code as Law Rebooted* (Edinburgh University Press 2022).

⁴⁹ Almquist and Lupton note that 'It is possible to employ the theory of affordances to support Latour's controversial notion that objects have agency, especially in situations when human (or animal) subjects interact with the object world in unexpected ways, beyond the designs of the designer. In such circumstances, the object takes on "a life of its own," becoming a new actant in an unpredictable situation or scenario.' Almquist and Lupton (n 38) 13.

frames of reference.⁵⁰ Verbeek provides a reconciliation and expansion of the analyses offered by Latour and Ihde. He foregrounds the way in which technologies mediate *both* action and perception and identifies the locus of mediation as the (already technologically mediated) *relations* between humans and their environment.⁵¹ If Ihde draws attention to how technologies operate as seats of influence, Verbeek emphasises the relationship between humans, technologies and their environment. Each maintains (in different but related ways) that technologies have a kind of agency; they may mediate action, perception, relations between humans and their environment.

Coeckelbergh offers a different but allied perspective which builds on narrative theory. He argues that technologies can and should be understood not just as ‘objects’ or ‘things’ but as narrators and meaning-makers. For Coeckelbergh, our narratives about technologies shape what they become,⁵² while, particularly in the case of AI, technologies ‘co-shape our narratives’ about ourselves and our world.⁵³ AI technologies, Coeckelbergh insists, are ‘hermeneutically active’; they are ‘interwoven with meanings and also change these meanings.’⁵⁴ Humans and technologies are both engaged in meaning-making; both become what they are in the process. In the case of legal technologies, we may expect that they may alter what it means to ‘do law’.

These various perspectives complement approaches rooted in Gibson’s theory of affordances. Our analyses in chapters 2 and 3 employ ‘affordances’ as a conceptual tool and make use of these complementary perspectives.

⁵⁰ Don Ihde, *Technology and the Lifeworld: From Garden to Earth* (Indiana University Press 1990).

⁵¹ Peter-Paul Verbeek, ‘Artifacts and Attachment: A Post-Script Philosophy of Mediation’, *Artifacts and Attachment: A Post-Script Philosophy of Mediation* (Amsterdam University Press 2005); Katinka Waelbers, ‘From Assigning to Designing Technological Agency’ (2009) 32 *Human Studies* 241.

⁵² Coeckelbergh’s recognition that technologies may become something other than what their developers intended recalls Ihde’s concept of the ‘multistability’ of technologies in their various contexts of use. Don Ihde, ‘Technology and Prognostic Predicaments’ (1999) 13 *AI & SOCIETY* 44.

⁵³ Mark Coeckelbergh, ‘Time Machines: Artificial Intelligence, Process, and Narrative’ (2021) 34 *Philosophy & Technology* 1623, 1630; Fleur Johns makes no reference to narrative theory or the hermeneutics of AI. However, her account of how digital interfaces transformed what it means to engage in humanitarian practices can be read as offering a concrete example of meaning-making by and through technology. Fleur Johns, *Digital Humanitarianism and the Remaking of International Order* (Oxford University Press 2023).

⁵⁴ Coeckelbergh (n 53) 1632.

2 The impact of data-driven technologies

Pauline McBride

In Chapter 1 we described a set of concepts and philosophical frameworks which are key to explaining why code-driven and data-driven technologies may transform law-as-we-know-it. Following Hildebrandt, we introduced the idea of law in its current *mode of existence* as an *affordance* of an information infrastructure which encompasses language, text and the printing press. We demonstrated that the concept of affordance, understood as the action possibilities arising from the relations between object and user, allows us to make sense of the dynamics of change brought about by new technologies. Complementary perspectives, which draw on postphenomenology, actor-network theory and narrative theory also shed light on the agential role of technologies. In this section we draw on these various perspectives to tease out the implications of data-driven technologies for law-as-we-know-it, the Rule of Law and the nature of the protection afforded by law.

2.1 The rise of data-driven legal technologies

Data-driven legal technologies – for our purposes those that employ machine learning techniques – are far from new. As early as 1974 Mackaay and Robillard used machine learning for the task of ‘prediction of judgment’.¹ Neural networks were applied to the prediction of judgment task in the field of construction litigation in the 1990’s.² Lex Machina, one of the first commercial organisations to use machine learning to assist lawyers to predict litigation outcomes, launched in 2010.³ Since 2015 there has been a conspicuous and sustained increase in research in the field of prediction of judgment, prompted in part by renewed interest in deep learning, the introduction and impact of transformer models and improved access to high quality, digitised data. Despite what Katz describes as ‘the march toward quantitative legal prediction’⁴ there are very few commercial products that offer prediction of judgment as a service.⁵ However, according to the European Commission for the Efficiency of Justice ‘public decision-makers are

¹ Ejan Mackaay and Pierre Robillard, ‘Predicting judicial decisions: The nearest neighbour rule and visual representation of case patterns’, *Predicting judicial decisions: The nearest neighbour rule and visual representation of case patterns* (De Gruyter 2020).

² David Arditi, Fatih E Oksay and Onur B Tokdemir, ‘Predicting the Outcome of Construction Litigation Using Neural Networks’ (1998) 13 *Computer-Aided Civil and Infrastructure Engineering* 75.

³ Lex Machina, ‘Lex Machina Celebrates 10 Years of Legal Analytics’ (*Lex Machina*) <<https://lexmachina.com/media/press/lex-machina-celebrates-10-years-of-legal-analytics/>> accessed 15 October 2023.

⁴ Daniel 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’ 62 *Emory Law Journal* 58, 912.

⁵ Case Crunch and CourtQuant, for example, are no longer trading, prompting Artificial Lawyer to ask, ‘... is litigation prediction dead?’ artificiallawyer, ‘Litigation Prediction Pioneer, CourtQuant, To Close’ (*Artificial Lawyer*, 7 October 2020) <<https://www.artificiallawyer.com/2020/10/07/litigation-prediction-pioneer-courtquant-to-close/>> accessed 15 October 2023.

beginning to be increasingly solicited by a private sector wishing to see these tools'.⁶ Senior UK judges have actively endorsed the use of prediction of judgment systems.⁷

Data driven technologies which carry out other tasks notably legal search, electronic discovery, document review and analytics, and compliance support have enjoyed commercial success.⁸ The online legal research service Westlaw has been using machine learning for case retrieval and natural language search for more than a decade.⁹ DiligenceEngine (now Kira), a contract analytics system that uses machine learning techniques, launched in 2010.¹⁰ Other commercial offerings (with launch dates in parenthesis) which employ or employed machine learning include LexPredict (2013), Ross Intelligence (2014), Luminance (2015), Predictrice, Jus Mundi (2019), Squirro (2019) and Della (2020), Manupatra and Afriwise.

Despite these developments Surden, writing in 2021, suggested that the use of machine learning in law is 'not extremely widespread'.¹¹ It is difficult to obtain a clear picture of the use of data-driven legal technologies. Surveys about use within the profession often have a low response rate. Data about investment in legal tech companies suggests year on year growth but typically does not distinguish between code-driven and data-driven systems and may be an unreliable indicator of use.¹² Use may vary from jurisdiction to jurisdiction. Large law firms likely made and continue to make use of a greater variety of data driven legal technologies than small firms. However, by 2019 use of AI-enabled technologies was sufficiently widespread for the American Bar Association to issue new regulations

⁶ European Commission for the Efficiency of Justice (CEPEJ), 'European Ethical Charter on the Use of Artificial Intelligence (AI) in Judicial Systems and Their Environment' <<https://www.coe.int/en/web/cepej/cepej-european-ethical-charter-on-the-use-of-artificial-intelligence-ai-in-judicial-systems-and-their-environment>> accessed 16 October 2023.

⁷ Scottish Legal News, 'Lord Chief Justice Anticipates AI Predictions of Case Outcomes' <<https://www.scottishlegal.com/articles/lord-chief-justice-anticipates-ai-predictions-of-case-outcomes>> accessed 8 August 2023; Sir Geoffrey Vos, 'Speech by the Master of the Rolls to the Bar Council of England and Wales' (*Courts and Tribunals Judiciary*, 18 July 2023) <<https://www.judiciary.uk/speech-by-the-master-of-the-rolls-to-the-bar-council-of-england-and-wales/>> accessed 16 October 2023.

⁸ Michael Mills lists legal research, e-discovery, compliance, contract analysis, case prediction and document automation as areas in which 'Artificial intelligence is hard at work in the law...' Michael Mills, 'Artificial Intelligence in Law: The State of Play 2016' (*Thomson Reuters Institute*, 23 February 2016) <<https://www.thomsonreuters.com/en-us/posts/legal/artificial-intelligence-in-law-the-state-of-play-2016/>> accessed 16 October 2023.

⁹ Noah Waisberg and Alexander Hudek, *AI for Lawyers: How Artificial Intelligence Is Adding Value, Amplifying Expertise, and Transforming Careers* (Wiley 2021).

¹⁰ David Curle and Steve Obenski, 'Ebook: AI-Driven Contract Analysis in Perspective and in Practice' (10 September 2020) <<https://kirasystems.com/forms/guides-studies/ai-driven-contract-analysis-perspective-and-practice/>> accessed 16 October 2023.

¹¹ Harry Surden, 'Machine Learning and Law: An Overview', *Research Handbook on Big Data Law* (Edward Elgar Publishing 2021) 179 <<https://www.elgaronline.com/display/edcoll/9781788972819/9781788972819.00014.xml>> accessed 18 August 2023.

¹² Chris Metinko, 'Legal Tech Makes Its Case With Venture Capitalists, Tops \$1B In Funding This Year' (*Crunchbase News*, 23 September 2021) <<https://news.crunchbase.com/venture/legal-tech-venture-investment/>> accessed 19 August 2023; Jane Croft, 'Why Are Investors Pouring Money into Legal Technology?' *Financial Times* (28 July 2022) <<https://www.ft.com/content/b6f0796e-0265-40c6-ad4c-a900cd788c39>> accessed 19 August 2023.

concerning the use of such technologies.¹³ If, nevertheless, in 2021 use of data driven legal technologies was not extremely widespread that may be about to change.

The launch of ChatGPT in 2022 may prove to be a watershed moment for the adoption of data driven legal technologies. The ‘generality and versatility of output’¹⁴ of so-called foundation models such as the GPT family make it particularly attractive to incorporate these models in commercial products. By February 2023, at least fourteen legal tech companies had announced that they were using GPT models in their product offerings.¹⁵ Casetext’s CoCounsel which is built on GPT-4 was launched in March 2023.¹⁶ By July of that year Casetext had been acquired by Thomson Reuters, the global publishing company and owners of Westlaw.¹⁷ Big law has also shown an interest in the capabilities of foundation models; Dentons,¹⁸ Allen and Overy¹⁹ and Troutman Pepper²⁰ have already launched systems built on OpenAI’s GPT family. A recent survey by Thomson Reuters found that 82% of respondents to a survey of mid-size and large firms in the US, Canada and the UK said they believe that ChatGPT and generative artificial intelligence (AI), can be readily applied to legal work, and 51% said that it should be.²¹

Data-driven legal technologies have also become increasingly sophisticated. Consider, for example, developments in commercial legal research systems. Many such systems offer conceptual search as

¹³ Lance Eliot, ‘Latest Insights About AI And The Law With A Keen Spotlight On The American Bar Association Remarkable Resolution 604’ (*Forbes*) <<https://www.forbes.com/sites/lanceeliot/2023/08/09/latest-insights-about-ai-and-the-law-with-a-keen-spotlight-on-the-american-bar-association-remarkable-resolution-604/>> accessed 19 August 2023. The text of the Resolution is available at <https://www.americanbar.org/content/dam/aba/directories/policy/annual-2019/112-annual-2019.pdf>.

¹⁴ Proposed Recital 60e in DRAFT Compromise Amendments on the Draft Report Proposal for a regulation of the European Parliament and of the Council on harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union Legislative Acts (COM(2021)0206 – C9 0146/2021 – 2021/0106(COD)).

¹⁵ Nicola Shaver, ‘The Use of Large Language Models in LegalTech’ (*Legaltech Hub*, 18 February 2023) <<https://www.legaltechnologyhub.com/contents/the-use-of-large-language-models-in-legaltech/>> accessed 23 August 2023.

¹⁶ Casetext, ‘Casetext Unveils CoCounsel, the Groundbreaking AI Legal Assistant Powered by OpenAI Technology’ <<https://www.prnewswire.com/news-releases/casetext-unveils-cocounsel-the-groundbreaking-ai-legal-assistant-powered-by-openai-technology-301759255.html>> accessed 22 April 2023.

¹⁷ Casetext, ‘Casetext to Join Thomson Reuters, Ushering in a New Era of Legal Technology Innovation’ (27 June 2023) <<https://casetext.com/blog/casetext-to-join-thomson-reuters-ushering-in-a-new-era-of-legal-technology-innovation/>> accessed 3 July 2023.

¹⁸ ‘Dentons to Launch Client Secure Version of ChatGPT’ <<https://www.dentons.com/en/about-dentons/news-events-and-awards/news/2023/august/dentons-to-launch-client-secure-version-of-chatgpt>> accessed 21 October 2023; ‘Product Walk Through: FleetAI, Dentons’ Gen AI Platform – Artificial Lawyer’ <<https://www.artificiallawyer.com/2023/10/09/product-walk-through-fleetai-dentons-gen-ai-platform/>> accessed 21 October 2023.

¹⁹ ‘A&O Announces Exclusive Launch Partnership with Harvey’ (*Allen Overy*, 15 February 2023) <<https://www.allenoverly.com/en-gb/global/news-and-insights/news/ao-announces-exclusive-launch-partnership-with-harvey>> accessed 15 October 2023.

²⁰ ‘Troutman Pepper Launches GPT-Powered AI Assistant’ (*Troutman Pepper - Troutman Pepper Launches GPT-Powered AI Assistant*, 22 August 2023) <<https://www.troutman.com/insights/troutman-pepper-launches-gpt-powered-ai-assistant.html>> accessed 24 August 2023.

²¹ Thomson Reuters, ‘New Report on ChatGPT & Generative AI in Law Firms Shows Opportunities Abound, Even as Concerns Persist’ (*Thomson Reuters Institute*, 17 April 2023) <<https://www.thomsonreuters.com/en-us/posts/technology/chatgpt-generative-ai-law-firms-2023/>> accessed 15 October 2023.

standard. Conceptual search enables lawyers to input natural language queries; the system finds and returns documents containing terms that are conceptually similar to the input terms.²² Contextual search is a more recent data-driven innovation. Lawyers can upload a document, such as a brief, into the system. The system assesses the context of the search from the document and uses the context to provide relevant results. One of the latest features offered by providers of legal research systems is the ability for the user to pose questions and receive answers. For example:

WestSearch Plus is a closed domain, non-factoid Question Answering system for legal questions that allows attorneys to zero in on the most salient points of law, related case law, and statutory law appropriate to their jurisdiction, in a way that traditional search and other legal research platforms cannot.²³

Instead of merely offering enhanced search functionality, the system provides responses that resemble legal advice. Casetext's CoCounsel will provide answers to research questions in the form of a memo, summarise documents including contracts or legal opinions, and prepare for a deposition.²⁴ ChatGPT (though clearly not marketed as a legal technology) can produce a draft contract (the jury is out on the utility of its outputs, even as a first draft).²⁵ As Ko notes, 'Increasingly, the output of artificially intelligent LegalTech resembles regulated activities that constitute legal practice.'²⁶ Against this background, Hildebrandt's anticipation of the emergence of data-driven 'law' appears perspicacious.

2.2 The affordances of data-driven technologies

Affordances, in Gibson's account, are 'subjective in that an actor is needed as a frame of reference.'²⁷ Most commercial data-driven legal technologies target lawyers, though judges and citizens might also interact with these systems. Yet the affordances of things are also 'objective in that their existence does not depend on value, meaning, or interpretation'.²⁸ Identifying 'objective' affordances represents a challenge, particularly where the context of use is a material *and* institutional environment such as law. Typically, our understanding of the action possibilities afforded by a commercial product is deeply

²² See, for example, Allison Hart, 'Elevate's Analyse Documents ELM Module: AI You Can Use' (*Elevate*, 13 May 2021) <<https://elevate.law/expertise/elevates-analyse-documents-elm-module-ai-you-can-use/>> accessed 5 November 2023.

²³ Tonya Custis and others, 'Westlaw Edge AI Features Demo: KeyCite Overruling Risk, Litigation Analytics, and WestSearch Plus', *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law* (Association for Computing Machinery 2019) <<https://doi.org/10.1145/3322640.3326739>> accessed 16 October 2023.

²⁴ Casetext (n 16). For a detailed analysis of CoCounsel see Pauline McBride and Masha Medvedeva, 'Casetext's CoCounsel through the Lens of the Typology' (*COHUBICOL*, 4 July 2023) <<https://www.cohubicol.com/blog/casetext-cocounsel-openai-typology/>> accessed 7 November 2023.

²⁵ Ken Adams, 'ChatGPT Won't Fix Contracts' (*Adams on Contract Drafting*, 9 December 2022) <<https://www.adamsdrafting.com/chatgpt-wont-fix-contracts/>> accessed 16 October 2023.

²⁶ Sebastian Ko, 'The Dark Side of Technology in Law: Avoiding the Pitfalls' in Susanne Chishti (ed), *The Legaltech Book: The Legal Technology Handbook for Investors, Entrepreneurs and FinTech Visionaries* (John Wiley & Sons 2020) 197.

²⁷ Joanna McGrenere and Wayne Ho, 'Affordances: Clarifying and Evolving a Concept' *Proceedings of Graphics Interface 2000 2* <<https://teaching.polishedsolid.com/spring2006/iti/read/affordances.pdf>>.

²⁸ *ibid.*

informed by the claims made by those who market the product. We may be inclined to interpret the action possibilities of the technology in relation to its intended user in the light of these claims, the target market and signifiers set out in the product interface. Identifying ‘objective’ affordances involves endeavouring to look beyond these framings²⁹ even while recognising that ‘our perception is always already mediated by language and interpretation’.³⁰

In this vein we suggest that machine-learning components in data-driven legal technologies offer the following broad affordances³¹ to users:³²

1. search of digitised materials using conceptual search (e.g. Westlaw Edge³³, Elevate’s Analyse Documents,³⁴ and Kira³⁵). Conceptual search allows users to obtain relevant results even when their input query does not contain words that appear in the information that is retrieved.³⁶
2. refining search results by providing information (in the form of documents) about the context of search (e.g. CARA AI,³⁷ Vincent³⁸)

²⁹ See Olia Lialina, ‘Once Again, the Doorknob: Affordance, Forgiveness, and Ambiguity in Human-Computer Interaction and Human-Robot Interaction’ *Media Theory* 60 (advocating for an approach to affordances which ‘allow[s] oneself and others to recognize [and, potentially, to act upon] opportunities and risks of a world that is no longer restrained to mechanical age conventions, assumptions, and design choices.’).

³⁰ Mireille Hildebrandt, ‘The Artificial Intelligence of European Union Law’ (2020) 21 *German Law Journal* 74, 76.

³¹ Davis notes that ‘theories of affordance have long been central to understanding and intervening in the development and analysis of technological systems, yet ML has remained outside of the design studies purview.’ Jenny L Davis, ‘“Affordances” for Machine Learning’, 2023 *ACM Conference on Fairness, Accountability, and Transparency* (ACM 2023) 330 <<https://dl.acm.org/doi/10.1145/3593013.3594000>> accessed 19 October 2023.

³² This analysis draws on our work in creating a Typology of Legal Technologies. L. Diver, P. McBride, M. Medvedeva, A. Banerjee, E. D’hondt, T. Duarte, D. Dushi, G. Gori, E. van den Hoven, P. Meessen, M. Hildebrandt, ‘Typology of Legal Technologies’ <<https://publications.cohubicol.com/typology/>>.

³³ ‘Westlaw Edge - A.I. Powered Legal Research’ <<https://legal.thomsonreuters.com/en/products/westlaw-edge>> accessed 30 October 2023.

³⁴ Allison Hart (n 22).

³⁵ Waisberg and Hudek (n 9) 136.

³⁶ Allison Hart (n 22).

³⁷ ‘What Is CARA A.I. and How Do I Use It? | Casetext Help Center’ <<https://help.casetext.com/en/articles/1971642-what-is-cara-a-i-and-how-do-i-use-it>> accessed 30 October 2023.

³⁸ Susan Cunningham, ‘Introducing Vincent: The First Intelligent Legal Research Assistant of Its Kind’ (*Medium*, 20 September 2018) <<https://blog.vlex.com/introducing-vincent-the-first-intelligent-legal-research-assistant-of-its-kind-bf14b00a3152>> accessed 30 October 2023.

3. obtaining insights (objectively, additional information) about collections of (usually textual) information (e.g. WestSearch Plus,³⁹ Lex Machina,⁴⁰ Uhura,⁴¹ Della,⁴² CoCounsel⁴³)
4. generating texts or textual responses (e.g. Mapping Bits,⁴⁴ CoCounsel⁴⁵)⁴⁶

All these affordances may change behaviours and produce real-world effects by allowing users to carry out certain kinds of actions. They are noteworthy because they depend on functionality which, at least in humans, requires language understanding and human reasoning. Data-driven technologies possess neither. Where, for example, data-driven legal technologies are used to draft contracts, make predictions, summarise case law, the affordances of the technologies are realised without the technologies engaging in (legal) reasoning.⁴⁷

2.2.1 Answering an objection

Some will object to the assertion that data-driven technologies do not engage in legal reasoning. They will point to the outputs of these systems. Look, they will say, GPT-4 passed the US Uniform Bar Exam;⁴⁸ prediction of judgment systems can achieve accuracy and F1-scores of over 90%;⁴⁹ large language models can be prompted to output predictions in the form of legal syllogisms,⁵⁰ a ‘chain of thought’⁵¹ or ‘reasoning steps’.

³⁹ ‘WestSearch Plus - Westlaw Edge’ <<https://legal.thomsonreuters.com/en/products/westlaw-edge/westsearch-plus>> accessed 30 October 2023.

⁴⁰ Marketing, ‘Legal Analytics by Lex Machina’ (*Lex Machina*) <<https://lexmachina.com/>> accessed 30 October 2023.

⁴¹ Uhura, ‘An Introduction to Information Extraction from Unstructured and Semi-Structured Documents’ (14 May 2021) <<https://uhurasolutions.com/2021/05/14/an-introduction-to-information-extraction-from-unstructured-and-semi-structured-documents/>> accessed 12 November 2023.

⁴² Della, ‘The Most Advanced AI on the Market for Legal Contract Review’ (*Della AI*) <<https://dellalegal.com/>> accessed 30 October 2023.

⁴³ Casetext (n 16).

⁴⁴ Dmitriy Skougarevskiy and Wolfgang Alschner, ‘Mapping Investment Treaties’ (*Mapping Investment Treaties*) <<http://mappinginvestmenttreaties.com/>> accessed 30 October 2023.

⁴⁵ Casetext (n 16).

⁴⁶ See Luciano Floridi and Massimo Chiriatti, ‘GPT-3: Its Nature, Scope, Limits, and Consequences’ (2020) 30 *Minds and Machines* 681 (suggesting that GPT-3 allows us to ‘mass produce good and cheap semantic artefacts’).

⁴⁷ Floridi and Chiriatti maintain that ‘The real point about AI is that we are increasingly decoupling the ability to solve a problem effectively—as regards the final goal—from any need to be intelligent to do so.’ *ibid.*

⁴⁸ Daniel Martin Katz and others, ‘GPT-4 Passes the Bar Exam’ (15 March 2023) <<https://papers.ssrn.com/abstract=4389233>> accessed 17 April 2023.

⁴⁹ Masha Medvedeva, *Identification, Categorisation and Forecasting of Court Decisions* (University of Groningen 2022) 48. Masha Medvedeva and Pauline McBride, ‘Legal Judgment Prediction: If You Are Going to Do It, Do It Right’, *In Proceedings of the Natural Legal Language Processing Workshop (NLLP’23), 2023 (forthcoming)*.

⁵⁰ Cong Jiang and Xiaolei Yang, ‘Legal Syllogism Prompting: Teaching Large Language Models for Legal Judgment Prediction’, *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law (Association for Computing Machinery 2023)* <<https://dl.acm.org/doi/10.1145/3594536.3595170>> accessed 24 October 2023.

⁵¹ Fangyi Yu, Lee Quartey and Frank Schilder, ‘Legal Prompting: Teaching a Language Model to Think Like a Lawyer’ (arXiv, 8 December 2022) <<http://arxiv.org/abs/2212.01326>> accessed 5 November 2023.

These claims, and their implications, deserve close scrutiny. Martínez points out various difficulties in verifying claims about GPT-4's performance in the bar exam.⁵² He also notes, for example, that GPT-4 performed rather less well overall in essay questions than in multiple choice questions. Prediction of judgment systems that employ an appropriate experimental set-up typically obtain rather more modest accuracy scores.⁵³ Jiang and Yang suggest that the fact that LLMs can be prompted to output text in the form of syllogisms indicates these systems are capable of deductive reasoning but accept that they are so far incapable of practical reasoning.⁵⁴ Yu et al. propose a method to prompt GPT-3 to 'think like a lawyer'.⁵⁵ They maintain that 'our analysis shows significant promise in prompt engineering for high-order LLM-based reasoning tasks' but concede 'it is questionable whether prompting actually teaches a LM to "think like a lawyer"'.⁵⁶ Thinking like a lawyer includes reasoning by analogy. Machine learning systems have shown poor performance on tasks which, for humans, require analogical reasoning.⁵⁷ Neither an output *in the form of* step-by-step reasoning nor an accurate output on a task which requires human reasoning, should be taken as evidence of *the exercise of reasoning*. Appearance is not the same as reality.⁵⁸

Data-driven technologies (for our purposes those that employ machine learning) from decision trees to GPT-4 employ statistical processes to learn patterns in their training data. A trained model takes an input and generates output (classifications, probability rankings, textual output) based on the patterns inferred from the data. A pre-trained large language model, for example, may take a textual prompt as an input and output text which is generated according to the model's 'statistical capacity to associate

⁵² Eric Martínez, 'Re-Evaluating GPT-4's Bar Exam Performance' (8 May 2023) <<https://papers.ssrn.com/abstract=4441311>> accessed 19 May 2023. See also Arvind Narayanan and Sayash Kapoor, 'GPT-4 and Professional Benchmarks: The Wrong Answer to the Wrong Question' (*AI Snake Oil*, 20 March 2023) <<https://aisnakeoil.substack.com/p/gpt-4-and-professional-benchmarks>> accessed 8 June 2023.

⁵³ Medvedeva (n 49); Medvedeva and McBride (n 49).

⁵⁴ Jiang and Yang (n 50).

⁵⁵ Yu, Quartey and Schilder (n 51).

⁵⁶ *ibid.* As Duarte points out, the legal syllogism merely provides a 'framework' for the presentation of legal arguments or justifications. The major and minor premises of the syllogism must first be constructed through a process of interpretation. Tatiana Duarte, 'Legal Reasoning and Interpretation' in Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt, 'Research Study on Text-Driven Law (Brussels 2023), Funded by the ERC Advanced Grant "Counting as a Human Being in the Era of Computational Law" (COHUBICOL) by the European Research Council (ERC) under the HORIZON2020 Excellence of Science Program ERC-2017-ADG No 788734 (2019-2024)' (*COHUBICOL*, 20 September 2023) 105, 106 <<https://www.cohubicol.com/news/publication-of-the-cohubicol-research-study-on-text-driven-law/>> accessed 15 October 2023.

⁵⁷ John Pavlus, 'The Computer Scientist Training AI to Think with Analogies' (*Scientific American*) <<https://www.scientificamerican.com/article/the-computer-scientist-training-ai-to-think-with-analogies/>> accessed 5 November 2023; Ian R Kerr and Carissima Mathen, 'Chief Justice John Roberts Is a Robot' (1 April 2014) 9 <<https://papers.ssrn.com/abstract=3395885>> accessed 5 November 2023.

⁵⁸ Daria Bylieva, 'Language of AI' <<https://soctech.spbstu.ru/en/article/2022.6.11/>> accessed 30 October 2023 (noting that Searle's 'Chinese room' experiment is relevant here). For a critique of the appearance/reality dichotomy see Mark Coeckelbergh and David J Gunkel, 'ChatGPT: Deconstructing the Debate and Moving It Forward' [2023] *AI & SOCIETY* <<https://doi.org/10.1007/s00146-023-01710-4>> accessed 23 August 2023.

words'.⁵⁹ The outputs can be impressive, but these models do not understand language as we do.⁶⁰ They have no conception of the world beyond their training data.⁶¹ They have no sense of overarching principles,⁶² legal or not; there is no hierarchy in training data. Machine learning systems can output text that resembles the *product* of legal reasoning, but the processes by which they output such text have nothing to do with the *exercise* of legal reasoning.⁶³ There is no poring over the constellation of facts at issue in a case or a contracting situation, no looking up the law, no exercise of judgment, no reflection on the demands of fundamental rights or of justice. There is no hesitation, no 'retracings and reattachment'⁶⁴ of speech acts and speakers,⁶⁵ no 'legal trajectory',⁶⁶ no possibility of satisfaction of the felicity conditions for the speech acts of law. Such systems are oblivious to law's 'regime of veridiction'.⁶⁷

⁵⁹ Floridi and Chiriatti (n 46).

⁶⁰ Gary Marcus and Ernest Davis, 'GPT-3, Bloviator: OpenAI's Language Generator Has No Idea What It's Talking about' (*MIT Technology Review*) <<https://www.technologyreview.com/2020/08/22/1007539/gpt3-openai-language-generator-artificial-intelligence-ai-opinion/>> accessed 29 October 2023; Emily M Bender and Alexander Koller, 'Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data' in Dan Jurafsky and others (eds), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (Association for Computational Linguistics 2020) <<https://aclanthology.org/2020.acl-main.463>> accessed 8 November 2023; Emily M Bender and others, 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜', *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (ACM 2021) <<https://dl.acm.org/doi/10.1145/3442188.3445922>> accessed 13 August 2023. Coeckelbergh and Gunkel agree that large language models 'manipulate signs without knowing that to which these tokens refer ... They generate different sequences of signs based not on actual meaning but according to statistically probable arrangements of difference.' However, for these authors this characteristic might be a feature rather than a bug provided one accepts a non-representational view of language. Coeckelbergh and Gunkel (n 58).

⁶¹ Bender and others (n 60).

⁶² Written legal norms must be interpreted in accordance with the sources of law and the principles of law relevant for the particular jurisdiction. Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56) 30.

⁶³ Benjamin Alarie, Anthony Niblett and Albert H Yoon, 'How Artificial Intelligence Will Affect the Practice of Law' (2018) 68 *The University of Toronto Law Journal* 106, 12 (noting that AI is not yet capable of 'reasoned judgment').

⁶⁴ Katja de Vries and Niels van Dijk, 'A Bump in the Road. Ruling Out Law from Technology' in Mireille Hildebrandt and Jeanne Gaakeer (eds), *Human Law and Computer Law: Comparative Perspectives* (Springer Netherlands 2013) 106.

⁶⁵ Bruno Latour, *The Making of Law: An Ethnography of the Conseil d'Etat* (Polity 2010) 218 ('Everything happens as if law were interested exclusively in the possibility of re-engaging the figures of enunciation by attributing to a speaker what he or she said. Linking an individual to a text through the process of qualification; attaching a statement to its enunciator by following the sequences of signatures; authenticating an act of writing; imputing a crime to the name of a human being; linking up texts and documents; tracing the course of statements: all law can be grasped as an obsessive effort to make enunciation assignable.').

⁶⁶ de Vries and van Dijk (n 64) 111, 113.

⁶⁷ In language that may be more familiar to those brought up on a diet of Anglo-American legal theory, such systems can have no 'internal point of view' about the bindingness of legal rules. HLA Hart, *The Concept of Law* (Third, Oxford University Press 2012) 115–117. Kerr and Mathen (n 57) 22, 27–30.

2.3 Data-driven technologies as change agents and influencers

‘What matter who’s speaking, someone said what matter who is speaking’⁶⁸

Samuel Beckett

2.3.1 Making way for a new normativity

Why should it matter that data-driven legal technologies are simultaneously capable of generating insights and texts and incapable of engaging in legal reasoning? In Samuel Beckett’s words, ‘What matter who’s speaking’⁶⁹ – and relatedly, why should it matter how they produce speech? On one view (we return to this in sections 2.3.2 and 2.3.3) it may not much matter, *so long as we do not imagine that such systems are speaking ‘legally’*, so long, that is, as we do not make the mistake of supposing that these systems are oriented to the felicity conditions of the speech acts of law. If we make that mistake,⁷⁰ we risk undermining or eroding law’s distinctive mode of ‘speaking’. Consciously or not, we open the door to a very different mode of existence of law. A shift in the register of what counts as ‘legal’, involving a move from ‘the realm of law to the realm of statistics’,⁷¹ implies a commensurate departure from law-as-we-know-it. In law, as Latour tells us, who speaks and how they speak matters.⁷²

We can speak ‘legally’ because law’s ‘enlanguaged’⁷³ mode of existence allows us to comprehend the retracings and attachments of law described by Latour and attribute legal effect. The very notion of legal effect presupposes the performative effect of a network of speech acts.⁷⁴ These performative effects establish and ‘define[...] the legal protection that is offered by modern positive law.’⁷⁵ Thus, the regime of enunciation of law allows us to orient our behaviour, anticipate legal outcomes, become legal subjects, engage with norms, speak of rights, recognise concepts such as ‘ownership’, ‘marriage’, ‘contracts’, ‘legal wrongs’. It allows us to make sense of the institutions of law, its connection to the state, our vulnerability to state-imposed sanctions. It is simultaneously a regime of enunciation and veridiction which allows us both to create new legal norms and to specify the conditions under which norms are

⁶⁸ Samuel Beckett, *Stories & Texts for Nothing* (Grove Press 1967) 85. Coeckelberg and Gunkel raise this question in relation to the outputs of large language models such as ChatGPT noting that ‘we now confront texts that have no identifiable author.’ Coeckelbergh and Gunkel (n 58).

⁶⁹ Samuel Beckett (n 68) 85.

⁷⁰ See Federico Cabitza, ‘A Reply: Lost in Communication? We Need a More Conscious and Interactive Use of AI’ (2022) 1 *Journal of Cross-disciplinary Research in Computational Law* <<https://journalcrcl.org/crcl/article/view/10>> accessed 10 November 2023 (noting that ‘Machines – especially those developed using Machine Learning [ML] techniques – can only make arguments and decisions, or even just ‘speak the truth’ [which cannot be contested], to the extent that we allow them.’).

⁷¹ Markku Suksi, ‘Formal, Procedural, and Material Requirements of the Rule of Law in the Context of Automated Decision-Making’ in Markku Suksi (ed), *The Rule of Law and Automated Decision-Making: Exploring Fundamentals of Algorithmic Governance* (Springer International Publishing 2023) 71. See also Laurence Diver and Pauline McBride, ‘Argument by Numbers: The Normative Impact of Statistical Legal Tech’ (2022) 3.

⁷² Law, Latour tells us, ‘insists on asking whether there is a path from one particular utterance to another, or between a given utterance and a given enunciator’. Bruno Latour, *An Inquiry into Modes of Existence: An Anthropology of the Moderns* (Harvard University Press 2013) 370.

⁷³ The term ‘enlanguaged’ was coined by Kiverstein and Rietveld. Julian Kiverstein and Erik Rietveld, ‘Scaling-up Skilled Intentionality to Linguistic Thought’ (2021) 198 *Synthese* 175.

⁷⁴ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56) 120.

⁷⁵ *ibid* 123.

'legal'. It ensures a high degree of coherence.⁷⁶ However imperfect,⁷⁷ this mode of existence of law preserves the rule of law value of respect for human autonomy.⁷⁸ As Hildebrandt notes,

Autonomy, accountability and justification all depend on prediction; we cannot act if we have no idea of the effects, we cannot be held accountable for what we could not have foreseen, and we cannot claim justification if we cannot not anticipate how others will evaluate our action.⁷⁹

If we make the mistake of supposing that data driven technologies can speak 'legally', we risk severing the connection between law and the shared communicative processes and understandings that make it possible for us to engage with law, to predict, foresee and anticipate legal effects.⁸⁰ We mangle the idea of what it means to engage in legal reasoning and interpretation, divorcing these practices from the 'web of meaning'⁸¹ and the iterative 'retracings and attachments' on which positive law relies. ⁸² We compound this risk if we imagine that these technologies generate legal norms, that is, enunciate speech acts which produce legal effects, in the same way as legislators or courts.⁸³

The concern is not for the mode of existence of law *as such*, but for how citizens and legal subjects experience and engage with the law. To the extent that we mistake or substitute the outputs of data-driven legal technologies for the rulings of judges, the advice given by lawyers, the views of citizens with

⁷⁶ A high, but not a perfect degree of coherence as Brownsword points out. Roger Brownsword and Karen Yeung (eds), *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes* (Hart 2008) 134–159.

⁷⁷ Mariano-Florentino Cuéllar, 'Cyberdelegation and the Administrative State' in Nicholas R Parrillo (ed), *Administrative Law from the Inside Out: Essays on Themes in the Work of Jerry L. Mashaw* (Cambridge University Press 2017) 156 (noting that 'Human deliberation is replete with all the limitations associated with human cognition, but implicit in the endeavor is an aspiration for dialogue and exchange of reasons that are capable of being understood, accepted, or rejected by policymakers, representatives of organized interests, and members of the public.').

⁷⁸ John Tasioulas, 'The Rule of Algorithm and the Rule of Law' (7 January 2023) <<https://papers.ssrn.com/abstract=4319969>> accessed 24 October 2023.

⁷⁹ Mireille Hildebrandt, *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology* (Paperback edition, Edward Elgar Publishing 2016) 59.

⁸⁰ Wiggins suggests that law in turn sustains shared communication between persons. David Wiggins, *Continuants: Their Activity, Their Being and Their Identity Twelve Essays* (Oxford university press 2016) 91 ('our sharing in a given specific animal nature and a law-sustained mode of activity is integral to the close attunement of person to person in language and integral to the human sensibilities that make interpretation possible.'). See also Tasioulas (n 78).

⁸¹ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56) 43, 50, 71.

⁸² We could of course choose to formally attribute legal effect to the output of these systems but it does not resolve the difficulty. de Vries and van Dijk draw the same conclusion about the implications of democratically sanctioned rule by the scripts of technology. de Vries and van Dijk (n 64) 119 ('We will then be in a situation in which every bit of script is created in accordance with a 'rule of law'. But when no legal acts of reattachments are enunciated these technological intermediaries will not partake in legal enunciation.').

⁸³ For the text output by a data driven technology to function as a speech act in law-as-we-know-it, more is required than that it should be intelligible or make sense as a set of words; it must 'make sense' by conforming to the felicity conditions of prior legal norms and setting the felicity conditions for future legal effects. In particular, Duarte notes that 'Savigny establishes two felicity conditions for interpretation: the interpreter must (i) attempt to reconstruct the intellectual trail of the legislator and (ii) acknowledge the historico-dogmatic whole of the legal system and perceive its relations with text.' Tatiana Duarte, 'Legal Reasoning and Interpretation' in Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56) 102.

some knowledge of the law, we make way for a very different kind of normativity than that of law-as-we-know-it and a different source and form of 'legal' effect.⁸⁴

2.3.2 Engines of influence

Let us suppose that we – citizens and lawyers – do not make the mistake of supposing that these systems can speak 'legally'. Let us assume that we remain cognisant of the very different 'reasoning' processes by which data-driven legal technologies generate outputs. Data-driven systems may nevertheless operate as 'engines of influence'⁸⁵ in their context of use.

This influence may be exerted in different ways. As Verbeek points out '[a]t the very moment human beings use them, artifacts change from mere "objects lying around" into artifacts-for-doing-something.'⁸⁶ In the case of data-driven legal technologies they become artifacts-for-search, -for-drafting-contracts, -for-prediction-of-judgment. They become situated in a practice or set of behaviours; their action possibilities are made manifest in use. They acquire meaning.⁸⁷

Coeckelbergh describes how data-driven technologies can be understood as 'shaping the narrative' of human actors, 'giving them roles', 'influencing meaning making' and 're-shaping a [...] practice'.⁸⁸ In the short term it may be that legal professionals will:

do less manual data assembly and initial analysis work but take on new tasks associated with interpreting and acting on the outputs of AI systems.⁸⁹

⁸⁴ We are therefore at odds with those who suggest, as Volokh does, that what matters is the output and not the method by which it is achieved. Eugene Volokh, 'Chief Justice Robots' 68 DUKE LAW JOURNAL.

⁸⁵ Pasquale and Cashwell use this phrase to describe prediction of judgment systems. Frank Pasquale and Glyn Cashwell, 'Prediction, Persuasion, and the Jurisprudence of Behaviorism' (8 November 2017) 3 <<https://papers.ssrn.com/abstract=3067737>> accessed 13 August 2023. Susskind, in a similar vein urges us to consider 'whether machines can deliver decisions at the standard of human judges or higher, not by replicating the way that judges think and reason but by using their own distinctive capabilities (brute processing power, vast amount of data, remarkable algorithms).' Richard E Susskind, *Online Courts and the Future of Justice* (First edition, Oxford University Press 2019). See also John Armour and Mari Sako, 'AI-Enabled Business Models in Legal Services: From Traditional Law Firms to next-Generation Law Companies?' (2020) 7 Journal of Professions and Organization 27. We are in good company. See, for example, Kerr and Mathen (n 57); Reuben Binns, 'Analogies and Disanalogies Between Machine-Driven and Human-Driven Legal Judgement' (2021) 1 Journal of Cross-disciplinary Research in Computational Law <<https://journalcrcl.org/crcl/article/view/5>> accessed 6 November 2023; Tasioulas (n 78).

⁸⁶ Peter-Paul Verbeek, *Moralizing Technology: Understanding and Designing the Morality of Things* (University of Chicago Press 2011) 97.

⁸⁷ Mark Coeckelbergh, 'The Grammars of AI: Towards a Structuralist and Transcendental Hermeneutics of Digital Technologies' [2022] Technology and Language, 3(2), 148-161 151 <<https://soctech.spbstu.ru/en/article/2022.7.9/>> accessed 30 October 2023.

⁸⁸ Mark Coeckelbergh, 'Time Machines: Artificial Intelligence, Process, and Narrative' (2021) 34 Philosophy & Technology 1623, 1627. Hildebrandt notes that 'Clark and Latour have pointed out that the usage of tools basically integrates them into our extended mind or delegates cognitive tasks to things that subsequently restrict or enlarge our 'action potential'.' Hildebrandt, *Smart Technologies and the End(s) of Law* (n 79) 108.

⁸⁹ James Faulconbridge, Atif Sarwar and Martin Spring, 'How Professionals Adapt to Artificial Intelligence: The Role of Intertwined Boundary Work' Journal of Management Studies 10

In the longer term, as Coeckelbergh says of ChatGPT, such systems may ‘change the way we think and experience the writing process and ourselves as writers.’⁹⁰

The concrete effects of the re-shaping of practice and of the meaning ascribed to data-driven legal technologies (that is, as for-doing-something) and their outputs may be hard to pin down.⁹¹ Sometimes, however, the influence of legal technologies and the potential effects of their use are more obvious. Ihde describes how technologies may present in relations of alterity, interacting with humans as a ‘quasi-other’.⁹² Data-driven legal technologies are often marketed as quasi-others – as ‘an automated associate assigned to write the first draft of your brief’ or a ‘CoCounsel’.⁹³ Uhura Solutions claim that their technology ‘reads and understands contracts just as humans do’.⁹⁴ Squirro say of their Augmented Intelligence Solutions that they ‘provid[e] a Smart Assistant-like experience’.⁹⁵

Bylieva, following Coeckelbergh, argues that language capability – or at least the ability to engage in dialogue – increases the likelihood of a technology being seen as a quasi-other.⁹⁶ This is relevant for data-driven systems which possess question answering functionality such as WestlawPlus,⁹⁷ Della⁹⁸ and

<<https://onlinelibrary.wiley.com/doi/abs/10.1111/joms.12936>> accessed 2 November 2023. Such a dynamic appears to be very much in play in Norkute et al.’s account of the experience of Thomson Reuters’ legal editorial team. Milda Norkute and others, ‘Towards Explainable AI: Assessing the Usefulness and Impact of Added Explainability Features in Legal Document Summarization’, *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Association for Computing Machinery 2021) <<https://dl.acm.org/doi/10.1145/3411763.3443441>> accessed 17 August 2023 (‘Since this AI model [a legal text summarisation system] has been in active use, the primary task of the editors has become to review and edit the machine-generated summaries rather than creating them from scratch based on the long input documents.’).

⁹⁰ Coeckelbergh and Gunkel (n 58).

⁹¹ Diver and McBride alert to the risk of ‘robotomorphy’ where we benchmark and align ourselves according to the standards of the technologies we create. See also Laurence Diver and Pauline McBride, ‘High Tech, Low Fidelity? Statistical Legal Tech and the Rule of Law’ [2022] *Verfassungsblog* <<https://verfassungsblog.de/roa-high-tech-low-fidelity/>> accessed 3 November 2023; Henrik Skaug Sætra, ‘Robotomorphy’ (2022) 2 *AI and Ethics* 5. Choi et al. claim to have conducted ‘the first randomized controlled trial of AI assistance’s effect on human legal analysis.’ Jonathan H Choi, Amy Monahan and Daniel Schwarcz, ‘Lawyering in the Age of Artificial Intelligence’ (7 November 2023) <<https://papers.ssrn.com/abstract=4626276>> accessed 13 November 2023. Empirical tests that address how the use of data-driven technologies can produce efficiencies are valuable, but do not take account of systemic effects.

⁹² Don Ihde, *Technology and the Lifeworld: From Garden to Earth* (Indiana University Press 1990) 97.

⁹³ Casetext (n 16).

⁹⁴ ‘Uhura Solutions | LinkedIn’ (27 May 2022) <<https://uk.linkedin.com/company/uhurasolutions>> accessed 31 October 2023.

⁹⁵ ‘Unlocking the Power of AI for Business Users’ (Squirro) <<https://squirro.com/why-squirro/>> accessed 31 October 2023.

⁹⁶ Bylieva (n 58) 121.

⁹⁷ ‘WestSearch Plus - Westlaw Edge’ (n 39).

⁹⁸ Artificial Lawyer notes that along with Della, ‘there are a number of legal AI companies that allow you to pose questions to the system and get answers back from a doc stack.’ artificiallawyer, ‘Meet Della AI – A New Challenger in the Doc Review/Analysis Market’ (*Artificial Lawyer*, 21 January 2020) <<https://www.artificiallawyer.com/2020/01/21/meet-della-ai-a-new-challenger-in-the-doc-review-analysis-market/>> accessed 13 July 2022.

Kira.⁹⁹ There is also anecdotal evidence to suggest that lawyers engage with some data-driven legal technologies as quasi-others.¹⁰⁰ This need not imply deference to the technology,¹⁰¹ but it points to its role as an engine of influence.¹⁰²

Moreover data-driven legal technologies employed in tasks such as legal research, document review, analysis and drafting implicitly or explicitly recommend, suggest, caution¹⁰³ and flag.¹⁰⁴ They may not dictate the content of advice, contracts, court documents or courses of action. However, they inevitably exert influence.¹⁰⁵ At the very least – as in the case of search, prediction of judgment, or systems used by court administrations to ‘triage’ cases to assess their relative importance – they influence a train of thought, a research strategy, consideration of options, courses of action.¹⁰⁶ A system which creates a first draft or reviews an earlier draft, is bound to have some effect on the final document.¹⁰⁷ In this way data-driven legal technologies exercise a degree of influence over the advice given by lawyers, the judgments

⁹⁹ Kira Systems launched question answering capability in 2020. ‘Kira Systems Launches Answers & Insights, A New-to-Market Capability in Contract & Document Analysis’ (24 August 2020) <<https://kirasystems.com/company-announcements/kira-systems-launches-answers-insights/>> accessed 31 October 2023.

¹⁰⁰ Waisberg and Hudek quote a client of Kira Systems as saying ‘... when I told executives at a client we were going to use Kira and explained what it was, the GC [General Counsel] said “I haven’t met “her” yet but I am glad we have her on the team.’ The authors themselves describe Kira as a ‘virtual Noah’. Waisberg and Hudek (n 9) 83 The names of some of these technologies lend themselves to a degree of anthropomorphism, including Kira, Della, Legal Robot. Anthony Niblett, a co-founder of Blue J Legal suggests that use of the system entails ‘... letting the data speak. It is not lawyers using their judgment about what is important.’ *ibid* 122.

¹⁰¹ Alicia Ryan’s comments are instructive: ‘You will get users with expectations at both ends of the spectrum. Either they think it’s [the AI system is] never going to work and they never give it a chance, or they think it’s AI and therefore it’s going to be perfect, so they just rely on it without checking.’ Alicia Ryan, ‘The ROI of AI: How a large firm determines it’ in Waisberg and Hudek (n 9) 141.

¹⁰² According to Romele, ‘Technologies, probably more than language, have their materialities and their affordances. And yet, they are also, or even mostly, signs of authority, intended to be believed and obeyed as they are.’ Alberto Romele, *Digital Habitus: A Critique of the Imaginaries of Artificial Intelligence* (Routledge 2024) 98.

¹⁰³ Wang’s review of the use of AI-powered systems in China’s judicial system notes that these ‘can “warn” human judges of similar cases and scenarios [where misjudgments or wrongful convictions were made] preventing the recurrence of past fallibilities’. Nu Wang, ‘“Black Box Justice”: Robot Judges and AI-Based Judgment Processes in China’s Court System’, *2020 IEEE International Symposium on Technology and Society (ISTAS)* (2020) 59 (citation omitted).

¹⁰⁴ Latour positions non-humans and humans, figurative and non-figurative technologies, flags and signs as actors. Wiebe E Bijker and John Law (eds), *Shaping Technology/Building Society: Studies in Sociotechnical Change* (Nachdr, MIT Press 2010) 244.

¹⁰⁵ Former Justice Mariano-Florentino Cuéllar notes that ‘people underappreciate the influence of certain technologies and information on their decisions.’ Cuéllar (n 77) 154.

¹⁰⁶ Some ‘prediction of judgment’ systems are explicitly presented as means of reducing ‘excessive variability’ in court decisions. European Commission for the Efficiency of Justice (CEPEJ) (n 6).

¹⁰⁷ Compose is one example of such a system. Compose claims that the system ‘cures “blank page” syndrome and starts attorneys off right’ ‘Better Briefs. Less Time. Fewer Headaches.’ (Compose) <<https://compose.law/>> accessed 31 October 2023; Shepherd offers an interesting reflection on the effects of use of AI in drafting and review. Jack Shepherd, ‘Lawyers: How Much Should You Rely on AI to Make First Drafts?’ <<https://jackwshepherd.medium.com/lawyers-how-much-should-you-rely-on-ai-to-make-first-drafts-69b7b0682c51>> accessed 1 November 2023.

issued by judges,¹⁰⁸ the content of contracts and ultimately the courses of action adopted by citizens.¹⁰⁹ By shaping the speech acts of judges, they affect the authoritative ascription of legal effect.¹¹⁰ By influencing lawyers' or citizens' expectations concerning the ascription of legal effect, they affect the outcomes to which legal effect is ascribed. In both cases – and despite their inability to 'speak legally' – they have an effect on legal effect.¹¹¹

2.3.3 New seats of power

We have cast data-driven legal technologies as agents of influence, considering, as it were, their prospective effect. However, it is important to recognise that they are also seats of power; they owe their existence to a network of actors with their own commitments, agendas, epistemologies, and regimes of veridiction. As Jongepier and Keymolen point out:

By focussing only on the output of a technology (the decision), we no longer take into account that this outcome is actually the interplay of a variety of associations of engineers, algorithms, data scientists, insurers, [...] experts, hardware, corporations, software, regulators and other stakeholders.¹¹²

We might add to that list: researchers, funding organisations, financial institutions, major accounting firms and, notably, legal publishers.¹¹³

¹⁰⁸ A senior judge in England and Wales has used ChatGPT to write part of a judgment. Hibaq Farah, 'Court of Appeal Judge Praises "Jolly Useful" ChatGPT after Asking It for Legal Summary' *The Guardian* (15 September 2023) <<https://www.theguardian.com/technology/2023/sep/15/court-of-appeal-judge-praises-jolly-useful-chatgpt-after-asking-it-for-legal-summary>> accessed 1 November 2023; Luke Taylor reports that 'A judge in Colombia has caused a stir by admitting he used the artificial intelligence tool ChatGPT when deciding whether an autistic child's insurance should cover all of the costs of his medical treatment.' Luke Taylor, 'Colombian Judge Says He Used ChatGPT in Ruling' *The Guardian* (3 February 2023) <<https://www.theguardian.com/technology/2023/feb/03/colombia-judge-chatgpt-ruling>> accessed 1 November 2023. Other judges are less impressed. 'Most Judges Haven't Tried ChatGPT, and They Aren't Impressed' (*The National Judicial College*) <<https://www.judges.org/news-and-info/most-judges-havent-tried-chatgpt-and-they-arent-impressed/>> accessed 8 November 2023.

¹⁰⁹ It is interesting to note that Volokh advocates that data-driven legal technologies engaged as AI judges should be assessed according to their persuasiveness rather than their accuracy. Volokh (n 84) 1152.

¹¹⁰ This may also be true of judges' clerks who may write the first drafts of judgments. Kerr and Mathen (n 57) fn 17 and associated text.

¹¹¹ For an in-depth discussion of legal effect and how technologies have effect on legal effect see Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56) 57–61, 134–137.

¹¹² Fleur Jongepier and Esther Keymolen, 'Explanation and Agency: Exploring the Normative-Epistemic Landscape of the "Right to Explanation"' (2022) 24 *Ethics and Information Technology* 49.

¹¹³ Davis notes that 'Three groups can be seen as predominating in the development of AI legal solutions.' These are legal publishers, the major accounting firms and 'venture capital supported entrepreneurs'. Anthony E Davis, 'The Future of Law Firms (and Lawyers) in the Age of Artificial Intelligence' (2020) 16 *Revista Direito GV* e1945, 10. Legal publishers have shown considerable interest in legal tech companies. Thomson Reuters acquired Casetext, SurePrep and ThoughtTrace. Wolters Kluwer acquired Della. Caroline Hill, 'What Wolters Kluwer's Acquisition of Della Means for Customers of Both Companies' (*Legal IT Insider*, 5 January 2023) <<https://legaltechnology.com/2023/01/05/what-wolters-kluwers-acquisition-of-della-means-for-customers-of-both-companies/>> accessed 4 November 2023. The reach of these publishing giants is considerable. In 2020,

How much influence will lawyers, judges and, for that matter, citizens exert over the design of data-driven legal technologies, the selection of training data,¹¹⁴ algorithms, experimental set-up, the metrics used in testing the systems, the choice of ‘explainability’ techniques (if any), the documentation of risks?¹¹⁵ These choices matter. They impact on outputs¹¹⁶ and affect the assessment of performance.¹¹⁷ They have a bearing on whether a technology will be adopted.¹¹⁸ Most importantly, they determine the affordances of the technology in its contexts of use. Design choices may make information about legal norms and the likely effects of those norms more or less accessible; they may reduce or increase the likelihood of the system being treated as a quasi-other and an authoritative source; they may facilitate or restrict human oversight and control and make it more or less easy to independently assess the outputs of the system.¹¹⁹ Ultimately the developers and providers of these systems have the power to determine what ‘law’ is communicated by their technologies, to whom, at what price and for which uses and purposes.

2.3.4 Summary

In this section, we demonstrated why and how data-driven legal technologies operate as change agents and influencers. They may have an effect on legal effect. As influencers, they may shape the speech acts of judges and affect the authoritative ascription of legal effect. They may shape lawyers’ or citizens’

Thomson Reuters announced that ‘Westlaw Edge is now in 100 per cent of U.S. law schools and nearly 50 per cent of AM Law 100 firms.’ The company also reported that it had signed ‘a multiyear contract ... with the administrative office of U.S. courts.’ Anita Balakrishnan, ‘All US Law Schools Now Use WestLaw Edge, Says Thomson Reuters’ (*Law Times*, 26 February 2020) <<https://www.lawtimesnews.com/resources/legal-technology/all-us-law-schools-now-use-westlaw-edge-says-thomson-reuters/326751>> accessed 25 October 2020.

¹¹⁴ Cantwell Smith notes the use of ‘vast collections of data sets, where we do not know what normative standards, registrations schemes, ethical stances, epistemological biases, social practices, and political interests have wrought their influence across the tapestry.’ Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgment* (The MIT Press 2019) 80.

¹¹⁵ Yeung points to the ‘chronic asymmetry of power between those who design, own, and implement these algorithmic decision-making systems and have access to the voluminous and valuable data upon which they rely, and the individuals whose lives they affect.’ Karen Yeung, ‘Why Worry about Decision-Making by Machine?’ in Karen Yeung and Martin Lodge (eds), *Algorithmic Regulation* (Oxford University Press 2019) 36 <<https://doi.org/10.1093/oso/9780198838494.003.0002>> accessed 3 November 2023. D’Ignazio and Klein point to the ‘privilege hazard’ associated with data science and artificial intelligence. Catherine D’Ignazio and Lauren F Klein, *Data Feminism* (The MIT Press 2020) 29.

¹¹⁶ Mart’s research into the variability of search results obtained through different legal research systems is instructive. Susan Mart, ‘Results May Vary’ [2018] ABA Journal <<https://scholar.law.colorado.edu/faculty-articles/964>>.

¹¹⁷ Medvedeva and McBride (n 49) (critiquing of the use of data contained in already-decided judgments for testing the performance of models used in the ‘prediction of judgment’ task); Cor Steging, Silja Renooij and Bart Verheij, ‘Taking the Law More Seriously by Investigating Design Choices in Machine Learning Prediction Research’, *Proceedings of the Sixth Workshop on Automated Semantic Analysis of Information in Legal Text (ASAIL 2023)*, June 23, 2023, Braga, Portugal (highlighting the effect of the choice of metrics for testing on performance scores).

¹¹⁸ Martínez argues that ‘To the extent that capabilities estimates for generative AI in the context [sic] law are overblown, this may lead both lawyers and non-lawyers to rely on generative AI tools when they otherwise wouldn’t and arguably shouldn’t ...’ Martínez (n 52) 3.

¹¹⁹ Passi and Vorvoreanu provide an insightful overview of the practices that may contribute to or militate against overreliance on AI generated outputs. They recommend that systems employ ‘cognitive forcing functions’ (ways of nudging people to reflect more carefully) and offer effective explanations of the system’s outputs to reduce the likelihood of overreliance. Samir Passi and Mihaela Vorvoreanu, ‘Overreliance on AI Literature Review’.

expectations concerning the ascription of legal effect and so influence the outcomes to which legal effect is ascribed. This is the here and now of law and legal practice.

However, use of these technologies may bring about more fundamental change. If we fail to distinguish between law's modes of enunciation and veridiction and the processes by which the technologies outputs texts, insights, answers, we put at risk the very mode of existence of law. We open the door to a very different kind and source of 'legal' effect, normativity and law. This is not yet the here and now of law and legal practice, but we should not be naïve. Financial pressures on justice systems and the interests of legal technology companies will play into the narrative that, at least for ordinary citizens and low value claims, data-driven 'law' is good enough.¹²⁰

2.4 A closer inspection of agentive effects: engines of influence, seats of power

Our examination in section 2.3 explored the dynamics through which data-driven legal technologies systems may impact on law and the practice of law. In this section we offer an evaluation of these effects by reference to Rule of Law values and the practices that sustain them.

2.4.1 Engines of influence, seats of power

That data-driven legal technologies can *shape* legal outcomes and norms created by judges is not news. Lawyers use these technologies in the hope of achieving better results. Few can imagine that such use is neutral in its effects. In many cases, the technologies make it possible to carry out analyses that would otherwise be impossible or prohibitively expensive.¹²¹ However, such use can also be problematic where there is overreliance on the technology and its outputs.¹²² Overreliance might be occasioned by laziness or poor practice, but it can also result from ignorance about the capabilities and limitations of the systems that are employed.¹²³ Mart's research about the very different results obtained by different commercial legal search systems is valuable; few will have appreciated the extent to which 'search results may vary'.¹²⁴ Similarly, Medvedeva's research concerning prediction of judgment systems – demonstrating that the high accuracy scores touted by prediction of judgment systems should not be taken at face value – provides a welcome reality check about the effectiveness and utility of prediction of judgment

¹²⁰ See Ashwin Telang, 'The Promise and Peril of AI Legal Services to Equalize Justice' (*Harvard Journal of Law & Technology*, 14 March 2023) <<https://jolt.law.harvard.edu/digest/the-promise-and-peril-of-ai-legal-services-to-equalize-justice>> accessed 12 November 2023.

¹²¹ For example, AI-powered contract review and analytics systems allow lawyers to carry out contract reviews at scale instead of reviewing a sample. Waisberg and Hudek (n 9) 134, 143.

¹²² The example of the New York lawyers who relied on ChatGPT is as instructive as it is notorious. Sara Merken, 'New York Lawyers Sanctioned for Using Fake ChatGPT Cases in Legal Brief' *Reuters* (26 June 2023) <<https://www.reuters.com/legal/new-york-lawyers-sanctioned-using-fake-chatgpt-cases-legal-brief-2023-06-22/>> accessed 7 November 2023.

¹²³ See Katherine Medianik, 'ARTIFICIALLY INTELLIGENT LAWYERS: UPDATING THE MODEL RULES OF PROFESSIONAL CONDUCT IN ACCORDANCE WITH THE NEW TECHNOLOGICAL ERA' 39 *CARDOZO LAW REVIEW* 1529 (suggesting that in the early days of adoption of e-discovery tools lawyers trusted these systems blindly.).

¹²⁴ Mart (n 116).

systems.¹²⁵ Few lawyers, we suspect, receive training about automation bias.¹²⁶ Systems may explicitly encourage reliance – even if their contract terms say something different.¹²⁷ Indeed a careful reading of the terms on which many data-driven legal technologies are supplied ought to put users of these systems on notice about their limitations!

Overreliance is a concern not only because of the risk of poor legal outcomes, but because it inappropriately puts power in the hands of the developers and providers of the technologies. However, at least in the case of lawyers, there are ways of managing the risk of overreliance – through monitoring use of the systems, training and education. Law schools, and legal regulatory bodies have a part to play here.¹²⁸ Whether as part of pre- or post-qualifying education lawyers should be equipped to be capable of understanding, in broad terms, the capabilities, limitations and likely effects of the systems they use.¹²⁹

The risk may also be tackled through system design; systems may be designed to prompt reflection and hesitation, employing what Passi and Vorvoreanu describe as ‘cognitive forcing functions’ (more prosaically, making you think).¹³⁰ Lawyers’ professional obligations of independence and competence should act as a buffer against overreliance provided that legal regulatory bodies do not give in to calls for relaxation of these standards of practice.¹³¹ Lawyers and citizens forget at their peril that an

¹²⁵ Medvedeva (n 49) 130–135. See also Medvedeva and McBride (n 49) (making the point also made by Pasquale and Cashwell, and Hildebrandt about the inappropriate nature of the experimental set-up of these systems and extending the analysis to more than 150 systems); Pasquale and Cashwell (n 85); Mireille Hildebrandt, ‘Algorithmic Regulation and the Rule of Law’ (2018) 376 *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 20170355; Mireille Hildebrandt, ‘Data-Driven Prediction of Judgment. Law’s New Mode of Existence?’ (2019) <<https://papers.ssrn.com/abstract=3548504>> accessed 16 August 2023.

¹²⁶ Noting the implications of automation bias Gentile argues that the legal profession will have to pass between ‘Scilla and Charybdis: the desire to “keep the law human” on the one hand, and blind faith in the “superior” powers of poorly understood and developing technologies (which will inevitably be flawed) on the other.’ Giulia Gentile, ‘LawGPT? How AI Is Reshaping the Legal Profession’ (*Impact of Social Sciences*, 8 June 2023) <<https://blogs.lse.ac.uk/impactofsocialsciences/2023/06/08/lawgpt-how-ai-is-reshaping-the-legal-profession/>> accessed 7 November 2023; For a discussion of automation bias and the mechanisms by which it operates see Kate Goddard, Abdul Roudsari and Jeremy C Wyatt, ‘Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators’ (2012) 19 *Journal of the American Medical Informatics Association: JAMIA* 121.

¹²⁷ See, in relation to Casetext’s CoCounsel Pauline McBride and Masha Medvedeva (n 24).

¹²⁸ Mireille Hildebrandt, ‘Grounding Computational ‘Law’ in Legal Education and Professional Legal Training’ in Bartosz Brożek, Oľia Kanevskaia and Przemysław Pałka (eds), *Research Handbook on Law and Technology* (Edward Elgar Publishing 2023) (forthcoming) (suggesting an approach which draws on the Typology of Legal Technologies). L. Diver, P. McBride, M. Medvedeva, A. Banerjee, E. D’hondt, T. Duarte, D. Dushi, G. Gori, E. van den Hoven, P. Meessen, M. Hildebrandt, (n 32).

¹²⁹ Mireille Hildebrandt, ‘Grounding Computational ‘Law’ in Legal Education and Professional Legal Training’ in Brożek, Kanevskaia and Pałka (n 128) (forthcoming).

¹³⁰ Passi and Vorvoreanu (n 119).

¹³¹ As to the risks presented by artificial intelligence tools to lawyers’ independence and competence see Peter Homoki, ‘Guide on the Use of Artificial Intelligence-Based Tools by Lawyers and Law Firms in the EU’ <https://www.ccbe.eu/fileadmin/speciality_distribution/public/documents/IT_LAW/ITL_Reports_studies/EN_ITL_20220331_Guide-AI4L.pdf>; Medianik (n 123) (calling for changes in the Model Rules of Professional Conduct to address the challenges of the use of AI tools by the profession).

independent judiciary and legal profession – with all that that entails – is crucial to democracy and the Rule of Law.¹³²

2.5 A closer inspection of agentic effects: making way for data-driven normativity

In section 2.3 we argued that if we fail to distinguish and keep a clear separation between law's modes of enunciation and veridiction and the processes by which the technologies outputs texts, insights and answers we open the door to a very different kind and source of 'legal' effect and normativity, and a different mode of existence of law. We make way for a form of data-driven normativity. This is a transformation of a different order.

On one view, formidable obstacles stand in the way of this vision of the future of law. One of these is the limited functionality of most current data-driven legal technologies. *In general*, these are not fact-finding machines, evidence gatherers or capable of near-simultaneous dialogue with multiple persons.¹³³ However, data-driven normativity need not depend on 'robot' judges; all that is necessary is that human judges or justice systems are *not merely influenced by but defer to* the outputs of data-driven technologies as though they spoke 'legally'. If, across societies and jurisdictions, we have not yet embraced this new order, we have certainly flirted with it here and there. In France, judges in the Courts of Appeal in Douai and Rennes conducted a three-month trial of AI-powered software designed to reduce variability in the rulings of judges.¹³⁴ The Shanghai intelligent assistive case-handling system for criminal cases (the '206 System') has a feature which can provide an alert as to whether (according to the analysis carried out by the system) a draft judgment deviates from the approach adopted in previous similar cases.¹³⁵ One of the most senior judges in the UK maintains that data-driven legal technologies 'may also, at some

¹³² As to the crucial role of an independent judiciary and legal profession for the rule of law see Margaret Satterthwaite, 'A/HRC/53/31: Reimagining Justice: Confronting Contemporary Challenges to the Independence of Judges and Lawyers - Report of the Special Rapporteur on the Independence of Judges and Lawyers' (OHCHR) <<https://www.ohchr.org/en/documents/thematic-reports/ahrc5331-reimagining-justice-confronting-contemporary-challenges>> accessed 7 November 2023 (noting that 'algorithmic decision-making brings promise and peril for the rule of law and for judicial independence').

¹³³ Note however that the claimed functionality of the Shanghai intelligent assistive case-handling system for criminal cases includes the ability to convert speech to text, provide summaries of evidence, associate evidence with claims, identify missing evidence, and produce a draft judgment. Yadong Cui, *Artificial Intelligence and Judicial Modernization* (Cao Yan and Liu Yan trs, Springer 2020) 158–163.

¹³⁴ European Commission for the Efficiency of Justice (CEPEJ) (n 6) 42.

¹³⁵ Cui (n 133) 158, 163; Nyu Wang and Michael Yuan Tian, "Intelligent Justice": AI Implementations in China's Legal Systems' in Ariane Hanemaayer (ed), *Artificial Intelligence and Its Discontents: Critiques from the Social Sciences and Humanities* (Springer International Publishing 2022) <https://doi.org/10.1007/978-3-030-88615-8_10> accessed 7 November 2023. Papagiannas notes the argument that 'by trying to achieve consistency through technology, the judicial system risks surrendering its power, shifting the nexus of decision-making power to the algorithms behind the smart systems.' Straton Papagiannas, 'Towards Smarter and Fairer Justice? A Review of the Chinese Scholarship on Building Smart Courts and Automating Justice' (2022) 51 *Journal of Current Chinese Affairs* 327.

stage, be used to take some (at first, very minor) decisions.¹³⁶ Yadong Cui, the former secretary and President of Party Committee of Shanghai Senior People's Court, strongly advocates the adoption of data-driven legal technologies describing the 'dream' of making 'justice a real science by combining justice and science and technology, using modern scientific and technological means.'¹³⁷

2.5.1 The texture of data-driven normativity

In the Research Study on Text-Driven Law we read that:

Because legal norms are enacted as written legal speech acts combined with the unwritten principles that are implied in the entirety of legal norms within a jurisdiction, their mode of existence is text-driven and thereby firmly grounded in natural language.¹³⁸

Moreover:

law is not a system of static rules where logical consistence is a goal in itself ... [it] is not a monologue based on deductive reasoning from immutable axioms, but a situated adversarial dialogue based on iterant *constructive re-interpretation* of the relevant legal norm. With law, we are not in the realm of mathematics but rather in the realm of practical reason, grounded in experience rather than logic.¹³⁹

The text-driven nature of law-as-we-know-it allows and obliges us to find a trade-off between certainty and uncertainty in law. It affords stability without stagnation. It underpins the extraordinary coherence and flexibility of law, makes it possible for us to participate in law as rational actors, to find in law both reasons for actions and justifications for decisions. It leaves room for contestation and ensures accountability.¹⁴⁰ What of the texture of data-driven normativity?

Legal theorists who have grappled with expressing the texture of *code-driven* normativity have offered evocative descriptions: computational legalism,¹⁴¹ or Double-Click justice,¹⁴² implying a mechanical application of the law,¹⁴³ 'not thinking about'¹⁴⁴ or hesitating over how rules may apply. Such lack of hesitation is also characteristic of data-driven normativity but there is an important distinction. Legal rules

¹³⁶ Sir Geoffrey Vos, 'Speech by the Master of the Rolls to the Law Society of Scotland' (*Courts and Tribunals Judiciary*, 14 June 2023) <<https://www.judiciary.uk/speech-by-the-master-of-the-rolls-to-the-law-society-of-scotland/>> accessed 12 November 2023.

¹³⁷ Cui (n 133) xix.

¹³⁸ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56) 32 (citation omitted).

¹³⁹ *ibid* 29 (original emphasis).

¹⁴⁰ *ibid* 1–5.

¹⁴¹ Laurence Diver, 'Computational Legalism and the Affordance of Delay in Law' (2021) 1 *Journal of Cross-disciplinary Research in Computational Law*.

¹⁴² Zenon Bankowski and Burkhard Schafer, 'Double-Click Justice: Legalism in the Computer Age' (2007) 1 *Legisprudence* 31.

¹⁴³ See *Lola v Skadden, Arps, Slate, Meagher & Flom, No 14-3845 (2d Cir 2015)* (reviewing authorities to the effect that the practice of law presupposes some exercise of judgement). See also Augustus Calabresi, 'Machine Lawyering and Artificial Attorneys: Conflicts in Legal Ethics with Complex Computer Algorithms' 34 *THE GEORGETOWN JOURNAL OF LEGAL ETHICS*.

¹⁴⁴ Bankowski and Schafer (n 142).

are explicitly (if imperfectly) represented in code-driven legal technologies.¹⁴⁵ In the case of data-driven normativity the connection to the rules or norms of law is much more attenuated. Machine learning systems that use decision trees may *learn* (technically, induce) rules from training data.¹⁴⁶ However, quite differently from code-driven systems, legal rules are not explicitly represented in data-driven systems. As Suksi points out:

While ... [the] previous decisions [forming part of the training data] may have a provision in the law as the point of departure, the new decision made based on a machine-learning algorithm *has the pool of previous decisions as the point of departure*, rather than the legal norm.¹⁴⁷

The outputs of these systems have no legal-normative inflection. In such outputs the normativity that informed the texts and other inputs used as training data is a vestigial trace.¹⁴⁸ This, if it is 'law', is a 'law' dissociated from legal normativity, the 'web' of legal powers,¹⁴⁹ the performativity of speech acts, the grounding in legal reasoning and interpretation.¹⁵⁰ Nevertheless, as we will show, the implications of this new normativity may very much depend on the extent to which the outputs of the technologies come to resemble outputs produced through the exercise of legal reasoning.

2.5.2 The implications of data-driven normativity: legal protection and the Rule of Law

We can address the implications of data-driven normativity by answering a provocation posed by Volokh. His provocation may be read as an answer to the second limb of our question – what matters who's speaking and *how they produce speech*. Volokh suggests that *how* AI-enabled technologies produce outputs matters not. According to Volokh, when we ask whether AI-enabled technologies are 'intelligent enough to do a certain task' it is the outputs that matter, not the methods by which the outputs are produced.¹⁵¹ Volokh extends this argument to the task of judging, advocating for the promotion of suitably trustworthy AI systems to the role of judge. He offers a vision of full-blown data-driven normativity.

Volokh's 'thought experiments' raise a crucial what-if question: what if these systems could issue judgments which, both in form and in substance, are indistinguishable from or would pass for those issued by human judges? Of course, this is a big 'what-if'. Such judgments, as Volokh acknowledges, would 'have to offer

¹⁴⁵ Diver (n 141).

¹⁴⁶ John Zeleznikow, 'The Benefits and Dangers of Using Machine Learning to Support Making Legal Predictions' (2023) 13 WIREs Data Mining and Knowledge Discovery e1505, 7, 8.

¹⁴⁷ Suksi (n 71) 72 (emphasis added).

¹⁴⁸ In the words of de Vries and van Dijk, law becomes 'pieces of historical evidence' de Vries and van Dijk (n 64) 119.

¹⁴⁹ Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56) 2.

¹⁵⁰ As to the relevance of all these factors for text-driven law, see Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56). It may be objected that Retrieval Augmented Generation (RAG) allows systems constructed on large language models to generate texts supported, for example, with links to legal texts. It is beyond the scope of this paper to address this point in full. However while there are advantages to RAG, its deployment imports known difficulties in information retrieval into the system. Parishad BehnamGhader, Santiago Miret and Siva Reddy, 'Can Retriever-Augmented Language Models Reason? The Blame Game Between the Retriever and the Language Model' (arXiv, 6 May 2023) <<http://arxiv.org/abs/2212.09146>> accessed 11 June 2023.

¹⁵¹ Volokh (n 84) 1137, 1138. For contrary views see Kerr and Mathen (n 57); Binns (n 85); Tasioulas (n 78).

explanatory opinions and not just bottom-line results.¹⁵² In law, a ‘result’ may be a finding of guilt or innocence, a sentence, the imposition of a fine or an order for damages, an order for divorce or any other order that may competently be granted by a court. ‘Explanatory opinions’ are justifications, provided by judges, which (1) are informed by and take the form of legal reasoning and (2) link the outcome of the case (the result) to the facts as established by the court (or agreed by the parties) and the relevant law.

As Schafer and Aitken point out:

It is essential that the legal process does not just try to give the right result. As a core requirement for the transparent administration of justice, the process has also to justify the result in a public way and to give reasons that can, at least in principle, be checked universally for correctness.¹⁵³

This is an aspect of legality and is deeply connected to the idea that laws should be made known to those affected by them.¹⁵⁴

Volokh’s commitment to the need for justifications as well as results explains why, of necessity, his argument is presented in a series of ‘thought experiments’. Given the current limitations of data-driven legal technologies we might expect to see – and do currently see – that such systems are incapable of producing legally relevant justifications.¹⁵⁵ Developers and providers of such systems can offer explanations of how the systems work. Such explanations contribute to transparency but shed no light on the justification for the output results.¹⁵⁶ Some systems can output information about the features of the input data which contribute most strongly to the system’s decision or classification.¹⁵⁷ This approach provides some information about why the system reached its decision, but these are explanations about the statistical significance of the features rather than justifications based on the norms of law. Hybrid systems which combine machine learning and traditional rule-based approaches might attempt, in effect, to retrofit a ‘justification’ derived through case-based argumentation onto a machine learning system

¹⁵² Volokh (n 84) 1138.

¹⁵³ Burkhard Schafer and Colin Aitken, ‘Inductive, Abductive and Probabilistic Reasoning’ in Giorgio Bongiovanni and others (eds), *Handbook of Legal Reasoning and Argumentation* (Springer Netherlands 2018) 310.

¹⁵⁴ Lon L Fuller, *The Morality of Law* (Rev ed, 15 print, Yale Univ Press 1978) 43, 49–51.

¹⁵⁵ In Gori’s words, ‘In asking an explanation of machine decisions, the meaning of the “why” and “because” which introduce, respectively, the question and the answer potentially belong to different linguistic games, each of which has its own vocabulary and forms of explanation, and make reference to different kinds of rules.’ Gianmarco Gori, ‘Law, Rules, Machines: “Artificial Legal Intelligence” and the “Artificial Reason and Judgment of the Law”’ (PhD Thesis, 2021).

¹⁵⁶ Elena Esposito, ‘Transparency versus Explanation: The Role of Ambiguity in Legal AI’ (2022) 1 *Journal of Cross-disciplinary Research in Computational Law* <<https://journalcrcl.org/crcl/article/view/10>> accessed 10 November 2023. For a judicial expression of this principle see *ECLI:NL:RBDHA:2020:1878*, *Rechtbank Den Haag*, C-09-550982-HA ZA 18-388 (English) [2020] Rb Den Haag ECLI:NL:RBDHA:2020:1878. Concerning the obligations on courts to give reasons see Ashley Deeks, ‘The Judicial Demand for Explainable Artificial Intelligence’ (2019) 119 *Columbia Law Review* 1829.

¹⁵⁷ Masha Medvedeva and others, ‘JURI SAYS: An Automatic Judgement Prediction System for the European Court of Human Rights’, *Legal Knowledge and Information Systems* (IOS Press 2020) <<https://ebooks.iospress.nl/doi/10.3233/FAIA200883>> accessed 10 November 2023.

output.¹⁵⁸ The ‘justification’ is not output by the machine learning system itself¹⁵⁹ and may be inconsistent with the output of that system.¹⁶⁰ Large language models can be prompted to output a prediction of judgment in the form of a ‘chain of thought’ or legal syllogism.¹⁶¹ However, such systems use the facts of already-decided cases for prediction; they are not reaching decisions about contested facts. The inability of data-driven legal technologies to link ‘results’ with legally relevant justifications remains a significant obstacle to the use of such technologies in legal decision-making.¹⁶²

Consider the implications of ‘bare’ results without justifications. There is no easy way for legal subjects to make sense of those results; no step-by-step reasoning, no way to identify the factual considerations which were judged to be relevant, no possibility of looking ‘backwards’ to the set of written legal norms which may have informed the result. It is, as Schafer and Aitken point out, impossible to check whether the result was justified according to legal norms. Moreover, while a ‘bare’ result may produce legal effects for the parties in the case, it has no wider legal normative effect. Figure 1, for example, shows the result (strictly, ‘order’ or ‘ruling’) in *Toivanen v Finland*¹⁶³:

FOR THESE REASONS, THE COURT,

1. *Declares*, unanimously, the application admissible;
2. *Holds*, by six votes to one, that there has been no violation of Article 6 of the Convention.

Figure 1

The order, per se, does not operate as a general legal norm. It can no more offer a guide to future conduct (or the likelihood of wanted or unwanted legal effects flowing from courses of action) for legal subjects than provide its own justification. The absence of a justification also has implications for contesting the order.¹⁶⁴ A bare order supplies no hint as to why the judge made the order. This presents difficulties

¹⁵⁸ Henry Prakken and Rosa Ratsma, ‘A Top-Level Model of Case-Based Argumentation for Explanation: Formalisation and Experiments’ (2022) 13 *Argument & computation* 159. Branting notes that ‘it seems very probable that useful decision support systems for explainable legal prediction must have a hybrid, two-stage design that permits explanation both in terms of legal predicates and in terms of factual features to span the gap between legal predicates and the language of ordinary discourse.’ Karl Branting, ‘Explanation in Hybrid, Two-Stage Models of Legal Prediction’, *XAILA@JURIX* (2020) 8 <<https://api.semanticscholar.org/CorpusID:235827410>>.

¹⁵⁹ Tasioulas suggests that ‘It is quite [sic] different thing, a fool’s gold version perhaps, to be given an ex post rationalisation of the decision that is causally inert, when the real cause of the decision is quite different.’ Tasioulas (n 78) 15.

¹⁶⁰ Prakken and Ratsma (n 158).

¹⁶¹ Jiang and Yang (n 50). As the authors candidly point out, the system, in its current form, is incapable of ‘interpret[ing] the law and reconstruct[ing] the facts.’

¹⁶² Zeleznikow suggests that ‘Perhaps, the most important challenge for using machine learning to support legal decision-making relates to explaining the derived decisions.’ Zeleznikow (n 146). Paradoxically, however, the ability for a system to output an explanation may increase the risk of overreliance. Cabitza (n 70).

¹⁶³ *Toivanen v Finland App no 46131/19 (ECtHR, 9 November 2023)*.

¹⁶⁴ Yeung (n 115) 24, 43 (pointing to ‘the legitimate interests of individuals in being able to identify a competent human person to whom they can appeal in contesting the decision’ and identifying ‘moral and legal rights to due process and participation, to be provided with an explanation of the reasons for adverse decisions, and to respect for one’s dignity and responsibility as a moral agent with capacity for self-reflection and self-control’). Article 6 of

not only for legal subjects who wish to contest the order but also for appellate courts who may wish to assess the soundness of the order by reference to the justification. There are implications therefore for judicial accountability; it is hard to hold a judge to account when it is impossible to scrutinise the basis on which their rulings have been made.¹⁶⁵ Orders without justifications can therefore be seen to be at odds with principles associated with the Rule of Law¹⁶⁶ including the idea that law should be publicly promulgated, general, prospective rather than retrospective, understandable, consistent, capable of being observed, stable and congruent.¹⁶⁷ Such orders are not conducive to the values of contestability, accountability and participation in the discourse of law – core values afforded (though not guaranteed) by law-as-we-know-it.

Between this and Volokh's 'trustworthy' systems there lies an entire spectrum of possibilities and problems which are far from solved. The inability of such systems to provide legally relevant justifications is not the only obstacle to the use of such technologies in legal decision-making. It is well-known that machine learning systems can replicate and amplify bias encoded in training data.¹⁶⁸ Bias can also be introduced through the design of machine learning systems and as a result of how information is presented by the system.¹⁶⁹ Other issues which affect performance of data-driven systems include an inability to generalise outside the distribution of the training data,¹⁷⁰ spurious correlations,¹⁷¹ model degradation,¹⁷² data

the European Convention on Human Rights (the right to a fair trial) implies a duty on the court to give reasons. Europäische Union, Europäischer Gerichtshof für Menschenrechte and Europarat (eds), *Handbook on European Law Relating to Access to Justice* (Publications Office of the European Union 2016); Desara Dushi, 'Human Rights in the Era of Automated Decision Making and Predictive Technologies' (*GlobalCampus of Human Rights - GCHR*, 11 April 2022) <<https://gchumanrights.org/gc-preparedness/preparedness-science-technology/article-detail/human-rights-in-the-era-of-automated-decision-making-and-predictive-technologies.html>> accessed 13 November 2023.

¹⁶⁵ Dushi points out that this in turn impacts on the right to an effective remedy provided by Article 13 of the European Convention on Human Rights. Desara Dushi (n 164).

¹⁶⁶ For a discussion of the Rule of Law see Gianmarco Gori, 'Rule of Law and Positive Law' in Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56).

¹⁶⁷ According to Fuller, these eight principles together express the 'internal morality of law'. Fuller (n 154) 38, 39, 41–90. For a discussion of the implications of Fuller's principles see Brownsword and Yeung (n 76) 118–128; Kristen Rundle, 'The Morality of the Rule of Law: Lon L. Fuller' in Jens Meierhenrich and Martin Loughlin (eds), *The Cambridge Companion to the Rule of Law* (Cambridge University Press 2021) 187 (arguing that Fuller's eight principles evince the 'distinctly moralized conception of reciprocity between lawgiver and legal subject that Fuller saw to be constitutive to the practice of the rule of law.').

¹⁶⁸ Ninareh Mehrabi and others, 'A Survey on Bias and Fairness in Machine Learning' (2021) 54 *ACM Comput. Surv.* <<https://doi.org/10.1145/3457607>>.

¹⁶⁹ *ibid.*

¹⁷⁰ Gary Marcus, 'The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence' (arXiv, 19 February 2020) <<http://arxiv.org/abs/2002.06177>> accessed 11 November 2023; Jiashuo Liu and others, 'Towards Out-Of-Distribution Generalization: A Survey' (arXiv, 27 July 2023) <<http://arxiv.org/abs/2108.13624>> accessed 11 November 2023.

¹⁷¹ Parikshit Bansal and Amit Sharma, 'Controlling Learned Effects to Reduce Spurious Correlations in Text Classifiers' in Anna Rogers, Jordan Boyd-Graber and Naoaki Okazaki (eds), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Association for Computational Linguistics 2023) <<https://aclanthology.org/2023.acl-long.127>> accessed 11 November 2023.

¹⁷² Daniel Vela and others, 'Temporal Quality Degradation in AI Models' (2022) 12 *Scientific Reports* 11654.

drift¹⁷³ and concept drift.¹⁷⁴ Retraining models can be costly but a failure to retrain may increase the risk of ‘freezing the future and scaling the past’.¹⁷⁵ It can prove difficult to allocate responsibility for failures or harms on account of the ‘many hands’ involved in design and use of the systems.¹⁷⁶ There is increasing recognition that such systems should be assessed not only for technical issues but a wide range of potential impacts and harms including environmental harms and dependency on harmful or exploitative labour conditions.¹⁷⁷

However, let us meet Volokh on his own terms. What if developers were capable of finding technical means of addressing interpretability and other issues relating to the performance and capabilities of these systems? What if the systems were capable of producing (in Volokh’s words) ‘trustworthy results’ accompanied by explanations in the form of legally relevant justifications? We must revisit the question, what matter who is speaking, and how they produce speech? It will not do, to answer *this* question, to say that we want judges to possess an ‘internal’ point of view on the normative effect of legal rules or that judges should demonstrate a commitment to law on the basis that these factors are relevant to the quality of the outputs.¹⁷⁸ We have already (for the purposes of argument only) conceded that the systems are technically capable of producing qualitatively acceptable results. The issue, if there is one, must lie elsewhere.

The key to the question, we suggest, involves considering the effect on (in Gibson’s language) the wider environment or niche. In other words, it is not clear that as Volokh suggests:

The normative question whether we ought to use AI judges should be seen as turning chiefly on the empirical question whether they reliably produce opinions that persuade the representatives that we have selected to evaluate those opinions.¹⁷⁹

For example, full-blown reliance on such systems would imply that humans would no longer judge. We would have abandoned the *practice* of the authoritative ascription of legal effect. Lawyers (if they still exist) would have to anticipate the outputs of the system. Just as ‘[t]o follow rules is to adopt a particular

¹⁷³ Grace A Lewis and others, ‘Augur: A Step towards Realistic Drift Detection in Production ML Systems’, *Proceedings of the 1st Workshop on Software Engineering for Responsible AI* (ACM 2022) <<https://dl.acm.org/doi/10.1145/3526073.3527590>> accessed 11 November 2023. As to the need to consider environmental dynamics associated with the deployment of machine learning technologies, see Sina Fazelpour, Zachary C Lipton and David Danks, ‘Algorithmic Fairness and the Situated Dynamics of Justice’ (2022) 52 *Canadian Journal of Philosophy* 44.

¹⁷⁴ Zhixue Zhao and others, ‘On the Impact of Temporal Concept Drift on Model Explanations’ in Yoav Goldberg, Zornitsa Kozareva and Yue Zhang (eds), *Findings of the Association for Computational Linguistics: EMNLP 2022* (Association for Computational Linguistics 2022) <<https://aclanthology.org/2022.findings-emnlp.298>> accessed 11 November 2023.

¹⁷⁵ Mireille Hildebrandt, ‘Code-driven Law: Freezing the Future and Scaling the Past’ in Simon Deakin and Christopher Markou (eds), *Is Law Computable?: Critical Perspectives on Law and Artificial Intelligence* (Hart Publishing 2020).

¹⁷⁶ Mark Coeckelbergh, *AI Ethics* (The MIT Press 2020) 113.

¹⁷⁷ Irene Solaiman and others, ‘Evaluating the Social Impact of Generative AI Systems in Systems and Society’ (arXiv, 12 June 2023) <<http://arxiv.org/abs/2306.05949>> accessed 21 October 2023.

¹⁷⁸ These being answers provided by Kerr and Mathen, and Tasioulas respectively. Kerr and Mathen (n 57); Tasioulas (n 78).

¹⁷⁹ Volokh (n 84) 1192.

form of life',¹⁸⁰ if we are to be able to speak legally, we must engage in the practice of speaking legally. If we are to be able to authoritatively ascribe legal effect, we must engage in the practice of authoritatively ascribing legal effect. We cannot evaluate outputs according to the standards of a practice if we do not engage in the practice.¹⁸¹ In the context of law, this dependency may be understood in terms of the

... circular normative relation [which] ties the conditions of felicity of legal speech acts, the texts in which such conditions are inscribed and the hermeneutic practices through which the latter "come to life".¹⁸²

As a result, in the event of wholesale substitution of such systems for human judges we will be incapable of assessing the trustworthiness of these systems according to human standards of judging.¹⁸³ We will no longer be able to hold the systems to account. We would also – inexplicably as it seems to us – have given up control over one of the mainstays of the system of checks and balances which provides protection against arbitrary power.

The current inability of data-driven legal technologies to produce legally meaningful justifications and (in a world where Volokh's vision reaches its apotheosis) our inability to evaluate the outputs of these technologies according to legal relevant standards of judging may be understood as two sides of the same coin. There is a gulf between the 'enlanguaged' normativity of text-driven law and the statistically mediated normativity of data-driven law. It is not clear that the gulf can be bridged. With law, as Latour maintains:

Either you are inside it and you understand what it does – without being able to explain it in another language – or you are outside it and you don't do anything 'legal'.¹⁸⁴

As it turns out, who speaks and how they produce speech matters.

2.6 Conclusion

Our exploration of the implications of data-driven legal technologies for law-as-we-know-it grounds that inquiry in the affordances of these technologies. Drawing on Gibson, we argue that the formidable affordances of such technologies are not matched by an ability on the part of such technologies to speak 'legally' and engage in legal reasoning and interpretation. We examine the agentic effects of the technologies, showing how they may act as engines of influence, shaping the outcomes to which legal

¹⁸⁰ Kerr and Mathen (n 57) 25.

¹⁸¹ Patterson, drawing on Wittgenstein observes that '... the meaning of a practice is an internal phenomenon. It is within the practice, and by virtue of the acts of the participants in the practice, that the practice has meaning. ... It is, therefore, against the specifics of a practice that claims for actions consistent with the practice are validated. Our perception of the objectivity of any particular decision is a function of the degree to which the act in question is in conformity with the demands of the practice as understood by the participants.' Dennis M Patterson, 'Law's Pragmatism: Law as Practice & Narrative' (1990) 76 *Virginia Law Review* 937.

¹⁸² Laurence Diver, Tatiana Duarte, Gianmarco Gori, Emilie van den Hoven and Mireille Hildebrandt (n 56) 48.

¹⁸³ This is perhaps in some way related to Volokh's suggestion that these systems be assessed for persuasiveness rather than correctness according to some standard of legal reasoning. Eugene Volokh, 'Chief Justice Robots' (2019) 68 *Duke Law Journal* 1135.

¹⁸⁴ Latour (n 72) 359.

effect is ascribed and the authoritative ascription of legal effect. We alert to their emergence as new seats of power. This, we suggest, is the here and now of law and legal practice. It is already clear that there are reasons for some concern about the implications of overreliance on the technologies. However, there is also reason to suppose that these concerns may be addressed through novel approaches to education and training, renewed emphasis on lawyers' professional duties of independence and competence, and a focus on the design of legal technologies.

Crucially we draw attention to a more profound risk: that by mistakenly treating the outputs of data-driven legal technologies as though the product of legal reasoning and interpretation, we make way for a new normativity and a different source and nature of legal effect. We explore the implications of a data-driven normativity marked by its dissociation from the 'web' of legal powers, the performativity of speech acts and the practices of legal reasoning and interpretation.

Taking Volokh's 'thought experiments' as a provocation we examine the implications of deferring – or delegating – to data-driven technologies in the context of judging. There are implications for law-as-we-know-it whether the outputs are trustworthy or not. In their current form, data-driven systems have a range of technical limitations which makes them unsuited for the task of judging. The most obvious limitation relates to their inability to output legally relevant justifications. This is presently an unsolved problem, but we should not assume either that it will remain so, or that such inability will prevent states or judicial systems from exploring their use. That inability has clear consequences for the Rule of Law values of contestability, accountability and participation. Even if the outputs are 'trustworthy' on Volokh's account, we imperil the values of accountability and participation through wholesale deference or delegation to data-driven legal technologies in the task of judging. The implications, in both scenarios, are much more far-reaching than market effects.

3 The impact of code-driven legal technologies

Laurence Diver

3.1 Introduction

For the purposes of this chapter, we adopt a definition of code-driven law as “legal norms or policies that have been articulated in computer code, either by a contracting party, law enforcement authorities, public administration or by a legislator.”¹ More specifically, we focus on the latter two categories, given the increasing interest and speed of development in the ‘rules as code’ space and the tangible efforts of public administrations in adopting it for real-world application.² This can be contrasted with so-called ‘cryptographic law’ based on blockchain applications, which despite huge interest (and hype) in recent years has nevertheless lost a great deal of public and scholarly attention in light of the ongoing collapse of initiatives based on cryptocurrency and non-fungible tokens (NFTs).³

Keeping the emphasis on the articulation of public legal norms in computer code, this chapter focuses primarily on ‘rules as code’ (RaC) as a subset of code-driven law. RaC initiatives are becoming very real, and the articulation in code of legal norms which lies at its heart speaks directly to the notion of computational law having a potential ‘effect on legal effect’.⁴ The position of our first Research Study is that attributed legal effect is a central mechanism by which law and the Rule of Law can provide protection, and so it is essential to enquire whether and how it might be impacted by the introduction of certain kinds of computation.

In the remainder of this chapter, we first consider what kinds of approach ‘rules as code’ refers to. We glean some definitions from prominent commentators, before attempting to make sense of the various classes of approach by placing them on a spectrum of the transformative impact on law that potentially they represent. We thread through analysis of potential impact at each point, before embarking in the

¹ Mireille Hildebrandt, ‘Code-Driven Law: Freezing the Future and Scaling the Past’ in Simon F Deakin and Christopher Markou (eds), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence* (Hart 2020) 67.

² For example, the Department for Work and Pensions in the UK recently contracted New Zealand-based firm Novallex to produce a rules as code implementation of Universal Credit, the UK’s first ‘digital-by-design’ benefit. See Sam Trendall, ‘DWP Looks to Embed Machine-Readable Laws into Digital “Universal Credit Navigator”’ (*PublicTechnology*, 24 October 2022) <<https://www.publictechnology.net/2022/10/24/economics-and-finance/dwp-looks-embed-machine-readable-laws-digital-universal-credit-navigator/>> accessed 20 July 2023.

³ Although private contracting does have undoubted relevance here, notably via the notion of ‘smart contracts’, these have proven to be much less transformative than proponents of blockchains have argued. Nevertheless, the notion of treating a smart ‘contract’ as a *legal* contract is closely connected with the risk of ‘conceptual slippage’ in the idea of legal effect that is central to the project (see Laurence Diver and others, ‘Research Study on Text-Driven Law’ (COHUBICOL 2023) 137 <<https://publications.cohubicol.com/assets/uploads/cohubicol-research-study-on-text-driven-law-final.pdf>> accessed 18 September 2023). The question of how legal effect is conceptualised and created is still relevant to the putatively private law of blockchain applications, however, and so the analysis here can be read with those in mind.

⁴ See ‘FAQ and methodology’ in Laurence Diver and others, ‘Typology of Legal Technologies’ (*Counting as a Human Being in the Era of Computational Law* (COHUBICOL), 2022) <<https://publications.cohubicol.com/typology/>> accessed 6 November 2023, and the discussion in section 3.5.4 below.

final section of the chapter on a broader consideration, from the COHUBICOL perspective, of the broader implications of code-driven law for legal protection and the Rule of Law.

Ultimately the analysis is nuanced: some rules as code approaches have the potential to enhance the practices that enable legal protection and the Rule of Law, while others reflect an instrumental idea of what a legal rule is and how it should be treated. The latter demonstrate a *legalistic* conception and application of the law, which is at odds with the idea of *legality*, whereby law is not just about rules, but also — and crucially — about affording spaces and procedures which allow for interpretation, deliberation and contestation. Those affordances, which are readily supported by text-driven ‘infrastructure’ of law-as-we-know-it, are what give law its capacity to provide protection, as well as democratic legitimacy to the processes by which the rules are produced. The interplay between RaC approaches and existing processes of government is complex and multifaceted. The hope in this chapter is to highlight where the introduction of computational methods can enhance the specifically and properly *legal* character of legislative rules, with all the procedural and interpretative connotations that ought to come with that term, while at the same time avoiding the potentially reductive impact of casting legal rules as no more than technical or commercial instructions for compliance.

3.2 Rules as Code

There is no settled consensus on what precisely RaC is, or what it should seek to be. This perhaps reflects the fact that people from various professional backgrounds are expressing an interest in its development and use, bringing quite varied conceptions of what it is and what it ought to be and do. In that sense we hope this Research Study, along with the prior study on text-driven law, might contribute something to the normative development of its scope and aims.

As Mohun and Roberts put it in the OECD’s working paper ‘Cracking the Code: Rulemaking for humans and machines’, ‘RaC envisions a fundamental transformation of the rulemaking process itself and of the application, interpretation, review and revision of the rules it generates.’⁵ They adopt de Sousa’s definition of RaC as ‘the process of drafting rules in legislation, regulation, and policy in machine-consumable languages (code) so they can be read and used by computers.’⁶

Kelly suggests that, ‘at its most basic’,

[RaC] is a granular agile project management methodology focussed on

- Creating a transparent algorithmic law representation, centred on decision tree diagramming and structured languages;

⁵ James Mohun and Alex Roberts, ‘Cracking the Code: Rulemaking for Humans and Machines’, vol 42 (2020) OECD Working Papers on Public Governance 42 16 <https://www.oecd-ilibrary.org/governance/cracking-the-code_3afe6ba5-en> accessed 23 October 2023.

⁶ Tim de Sousa, ‘Introduction: What Is Rules as Code?’ (*Rules as Code Handbook*, 19 March 2019) <<https://github.com/Rules-as-Code-League/RaC-Handbook/wiki/1-Introduction:-What-is-Rules-as-Code%3F>> accessed 6 November 2023.

- Secure, cloud-based production platforms, allowing iteration, testing, access, and maintenance.⁷

Waddington highlights a variety of approaches that come under the RaC umbrella. He suggests that the approach is

not wedded to any ... technology solutions so much as to the idea that the “coding” (or mark-up) of the legislation should be widely usable, traceable to the legislation, rather than adding material to reflect assumptions about procedures or implementation.⁸

He makes a nuanced normative argument about what RaC should (and should not) seek to achieve:

what if, hype aside, Rules as Code is not really intended to be magically transformative, and is not really about automating legal decisions or about programs that implement law themselves? ... [Rules as Code] is about applying human intelligence, rather than AI, and about the less glamorous ways in which computers are already handling law and could do better in aiding humans.⁹

This is echoed in visions for RaC that are as much about the process of developing rules as the means by which they are represented, processed and enforced. Casanovas, for example, suggests that RaC is

not a new technology (there is no way to present it as such) but an attitude that includes technological and political planning for policy making and a clear will to cope with the demands of the digital age.¹⁰

In this context of policy development, de Sousa and Andrews suggest that by making laws ‘machine consumable, we can take a whole new approach to testing them, and modelling different legislative approaches’.¹¹ To achieve this will require new ways of drafting legislation that mean we ‘draft the code and the human-readable text at the same time, and allow them to influence each other’.¹²

Once the policy has been developed and the RaC translations created (whether directly by the legislature/executive, or subsequently by third parties), the focus can shift to enforcement, compliance and the provision of legal advice. Extensible platforms such as DataLex can ‘be used to develop legal

⁷ Adrian Kelly, ‘Evolution of Digital Law’ [2023] *The Loophole - Journal of the Commonwealth Association of Legislative Counsel* 43, 46. Kelly is a legislative drafter and is active in various governmental RaC projects around the world.

⁸ Cf. Matthew Waddington, ‘Research Note: Rules as Code’ (2020) 37 *Law in Context. A Socio-legal Journal* 179, 180. Waddington also a legislative drafter active in the RaC community.

⁹ *ibid* 182.

¹⁰ Pompeu Casanovas, ‘Comments on Cracking the Code. A Short Note on the OECD Working Paper Draft on Rules as Code’, *Comments on Cracking The Code: Rulemaking For Humans And Machines (August 2020 draft)* (LawTech La Trobe Research Group 2020) 19 <<https://www.latrobe.edu.au/law/research/la-trobe-law-tech>> accessed 23 February 2022.

¹¹ Tim de Sousa and Pia Andrews, ‘When We Code the Rules on Which Our Society Runs, We Can Create Better Results and New Opportunities for the Public and Regulators, and Companies Looking to Make Compliance Easier’ (*The Mandarin*, 30 September 2019) <<https://www.themandarin.com.au/116681-when-machines-are-coding-the-rules-on-which-our-society-runs-we-get-better-results-new-opportunities-for-the-public-and-regulators-and-companies-looking-to-make-compliance-easier/>> accessed 6 November 2023.

¹² *ibid*.

reasoning applications in areas such as legal advisory services, regulatory compliance, decision support and Rules as Code'.¹³ 'Low-code' and 'no-code' platforms like Blawx and Neota enable the declaration of rules using intuitive graphical user interfaces that hide the logical rule engines beneath.¹⁴ Domain experts who might not have experience with logic or general-purpose programming can thus be closely involved in the process of defining rules.¹⁵

The impact of RaC will depend very much on who is afforded what uses by the system, and at what stage of the 'lifecycle' of a legal norm. The 'who' can be anyone from legislative drafters to public administrators to commercial enterprises to citizens to judges. The point of application can be anything from point of initial conception to the developing and passing of a norm into law to the interpretation of its terms by those subject to them on to the authoritative determination of its meaning by a court (if that ever happens) — as well as at many points in between.

3.2.1 The COHUBICOL lens

Taking a step back is part of COHUBICOL's approach: attempting to tease out the implicit and explicit assumptions that are reflected in the design of legal technologies and the contexts where they are intended to be deployed. It is essential to properly delineate the ways in which humans (both citizen 'users' and legal practitioners, of all kinds) might be aided by technologies like RaC, to ensure that core legal values can be preserved and, where possible, enhanced by the appropriate use of such technology.

3.2.1.1 Legality and legalism

Casting law as 'regulation', and citizens as 'consumers of rules' or 'rule-takers', risks adopting a reductive and technocratic view of legal norms as merely instruments of policy, and citizens and other legal subjects as passive targets whose duty is simply to comply. This flattens the relationships that properly constitute *legality*, as opposed to *legalism*, in which there is a "reciprocity of expectations" between legislator and

¹³ 'DataLex: AustLII's Legal Reasoning Application Platform' <<https://www.datalex.org/>> accessed 6 November 2023. For the COHUBICOL analysis of DataLex through the lens of the Typology of Legal Technologies, see 'DataLex' in Diver and others (n 4).

¹⁴ See 'Blawx.Com – User Friendly Rules as Code' <<https://www.blawx.com/>> accessed 7 November 2023 and, for the COHUBICOL analysis of Blawx through the lens of the Typology of Legal Technologies, see 'Blawx' in Diver and others (n 4). On Neota, see 'Automating Processes Just Got Easier' (*Neota*) <<https://neota.com/>> accessed 7 November 2023.

¹⁵ On low- and no-code approaches, see Jason Morris, 'Code vs. No-Code' (*Rules as Code Diary*, 24 February 2022) <https://JasonMorrisSC.github.io/post/2022-02-24_no-code/> accessed 7 November 2023.

citizen.¹⁶ Legality is about the reciprocal rule of law, whereas legalism is about top-down rule by law.¹⁷ A failure fully to uphold this reciprocity in a single individual instance might not matter too much to those involved — true legality is an aspiration, and unlikely to be attained in every case. But in aggregate, the failure to do so could threaten the social, civic and professional ‘anchoring practices’ of ‘interpretation, justification, contestation and creative action’ that both give rules their legal character and enable the law to afford protection.¹⁸

These practices are forever in tension, requiring constant reinvigoration for law to have legitimacy and effectiveness.¹⁹ When the (technological) conditions that enable those practices change, the question that must be asked is whether or and how the practices themselves will change, and possibly falter, possibly in unforeseen or non-obvious ways. That is the essential question the COHUBICOL project seeks to grapple with.

From that perspective, it is essential not to assume any determinism in the role played by technologies involved in the law: we cannot presume that they will have only beneficial or only negative effects, if even we can assume their introduction will have some kind of impact.²⁰ Instead, to properly identify that impact we need to anticipate (i) who will be afforded what capacities by them and in what circumstances, (ii) which existing affordances will be changed or removed, and (iii) what impact might they have on the conceptual underpinnings of the law and its particular ‘mode of existence’ (these three elements are of

¹⁶ Cf. Lon L Fuller, *The Morality of Law* (Yale University Press 1977) 209. (“With a legal system... the existence of a relatively stable reciprocity of expectations between lawgiver and subject is *part of the very idea of a functioning legal order.*” (my emphasis)). There is a significant literature on the crucial differences between legalism and legality; the former viewing law as chiefly concerned with top-down application of rules, and the latter having a more reflexive quality that respects interpretation, autonomy and judgement. See, for example, Mireille Hildebrandt, ‘Radbruch’s Rechtsstaat and Schmitt’s Legal Order: Legalism, Legality, and the Institution of Law’ (2015) 2 *Critical Analysis of Law* 42; Jeremy Waldron, ‘The Rule of Law and the Importance of Procedure’ (2011) 50 *Nomos* 3; Zenon Bańkowski and Neil MacCormick, ‘Legality without Legalism’ in Werner Krawietz and others (eds), *The Reasonable as Rational? On Legal Argumentation and Justification; Festschrift for Aulis Aarnio* (Duncker & Humblot 2000); Zenon Bańkowski, ‘Don’t Think About It: Legalism and Legality’ in Mikael M Karlsson, Ólafur Páll Jónsson and Eyja Margrét Brynjarsdóttir (eds), *Rechtstheorie: Zeitschrift für Logik, Methodenlehre, Kybernetik und Soziologie des Rechts* (Duncker & Humblot 1993). For discussion of legalism specifically in relation to computation, see Laurence Diver, *Digisprudence: Code as Law Rebooted* (Edinburgh University Press 2022) ch 3; Zenon Bańkowski and Burkhard Schafer, ‘Double-Click Justice: Legalism in the Computer Age’ (2007) 1 *Legisprudence* 31; Roger Brownsword, ‘Technological Management and the Rule of Law’ (2016) 8 *Law, Innovation and Technology* 100; Philip Leith, ‘The Application of AI to Law’ (1988) 2 *AI & Society* 31.

¹⁷ Mireille Hildebrandt, ‘Legal Protection by Design: Objections and Refutations’ (2011) 5 *Legisprudence* 223, 234 (‘This means that individual citizens have a means to challenge the administration’s interpretation of enacted law, thus preventing a mere rule by law that employs the law as a neutral instrument to achieve the goals of policy makers’). When the rules in question are embedded in digital architectures rather than text, this ‘rule by code’ threatens to reach a peak: ‘computational legalism’. See Diver, *Digisprudence* (n 16) pt 1 (‘Computational Legalism and the Rule(s) of Code’).

¹⁸ Stephen J Toope, *A Rule of Law for Our New Age of Anxiety* (Cambridge University Press 2023) 173.

¹⁹ Cf. Radbruch’s antinomian concept of law: Gustav Radbruch, ‘Legal Philosophy’ in Kurt Wilk (ed), *The Legal Philosophies of Lask, Radbruch, and Dabin* (Harvard University Press 1950); Hildebrandt, ‘Radbruch’s Rechtsstaat and Schmitt’s Legal Order: Legalism, Legality, and the Institution of Law’ (n 16).

²⁰ Reflecting Kranzberg’s maxim ‘technology is neither good nor bad; nor is it neutral’ (Melvin Kranzberg, ‘Technology and History: “Kranzberg’s Laws”’ (1986) 27 *Technology and Culture* 544, 545).

course deeply intertwined).²¹ Ultimately the hope is that such technologies can be embraced where they facilitate and strengthen the specific type of protection that law affords, and resisted where they promote interests or centres of power that could or would otherwise have been constrained by the Rule of Law.

3.3 The texture of code-driven normativity

We mentioned above the idea that one way to make sense of the RaC landscape is to consider the spectrum of its potential impact on the law. More specifically, we want to tease out the difference between what a legal rule does and what a code rule does, and where particular approaches to RaC sit in relation to those two very different types of ‘normativity’ — the ways in which behaviour, action and practice are shaped through inducement, enforcement, inhibition or prohibition.²²

In the first Research Study, we considered the texture or fabric of text-driven normativity from various angles, trying to identify its qualities and the conditions that makes it possible. One task in this chapter is to attempt to do something similar for code-driven law (specifically RaC).

Different technologies exert different amounts of ‘normative force’, from suggesting behaviour and action, through to guiding them in ways that can be resisted, on to defining their character and limits from the outset, with no possibility of reconfiguration or resistance.²³ From a Science, Technology and Society studies perspective this is true of all technologies; they inevitably shape the practices they are embedded within. This shaping is often imperceptible and can be a constitutive part of the practice, sometimes in ways that fade into the background.²⁴ Text is perhaps an example of such a technology.²⁵ Despite it playing a fundamental role in shaping the nature of legal rules and the character of their application, this fact can be easily missed because it is so familiar to us.²⁶

²¹ On the law’s current mode of existence, see Mireille Hildebrandt, ‘1. Introduction: The Mode of Existence of Text-Driven Positive Law’, *Research Study on Text-Driven Law* (COHUBICOL 2023) <<https://publications.cohubicol.com/research-studies/text-driven-law/chapter-3/legal-reasoning-and-interpretation/legal-reasoning-and-interpretation/>> accessed 18 September 2023.

²² The distinction is set out in detail in Mireille Hildebrandt, ‘Legal and Technological Normativity: More (and Less) than Twin Sisters’ (2008) 12 *Techné: Research in Philosophy and Technology* 169. This chimes with Davis and Chouinard’s normative framing of affordances in terms of whether they request, demand, allow, encourage, discourage, or refuse a particular behaviour or action. See Jenny L Davis and James B Chouinard, ‘Theorizing Affordances: From Request to Refuse’ [2017] *Bulletin of Science, Technology & Society*.

²³ Diver, *Digisprudence* (n 16) ch 2 (‘Code is more than law: a design perspective’).

²⁴ See e.g. Katja de Vries and Niels van Dijk, ‘A Bump in the Road. Ruling Out Law from Technology’ in Mireille Hildebrandt and Jeanne Gaakeer (eds), *Human Law and Computer Law: Comparative Perspectives* (Springer Netherlands 2013) (discussing the salience of the media that underpin legality, in light of the ‘practice turn’ in law).

²⁵ On text and the printing press as transformative technologies, see Walter J Ong, *Orality and Literacy: The Technologizing of the Word* (3rd edn, Routledge 2012); Elizabeth L Eisenstein, *The Printing Revolution in Early Modern Europe* (2nd edn, Cambridge University Press 2012).

²⁶ On text as a technology underpinning text-driven law, see Mireille Hildebrandt, ‘2.4 The Texture of Modern Positive Law’ in Diver and others (n 3). See also Mireille Hildebrandt, *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology* (Edward Elgar Publishing 2015); Mireille Hildebrandt, ‘A Vision of Ambient Law’ in Roger Brownsword and Karen Yeung (eds), *Regulating Technologies: legal futures, regulatory frames*

Crucially, the normative force of technology operates both through the technology in which the rules are embedded, be that text or code, and on the people involved throughout the lifecycle of a rule — from drafters and policymakers to public administrators, compliance officers, citizens, litigants and judges. While it might be tempting to think of legal rules (and indeed any rule) as cleanly logical and susceptible to application free of external influence, their interpretation and the technological means by which they are produced, accessed and enforced all have an impact on their real-world operation.

It follows that when it comes to RaC, the normative impact comes not only from the RaC-modelled rules themselves (in whatever specific form they take), but also the technologies used to produce, disseminate and enforce them.²⁷ On the ex-ante production side of rules, we can think about how RaC technologies represent or facilitate foundational concepts of law, e.g. policy, rights, personhood, legal effect and justice.²⁸ On the ex-post side, where the rules then play some role in real-world contexts, we can consider whether RaC affects the interpretative, adjudicative and determinative processes of law. In either case, the potential impacts will be different for different actors, who will include citizens, litigants, practitioners, public administrators and judges. The interests and role of each are important when considering what it is that legal rules are meant to do, and how.

Text-driven law has a very specific type of normativity. Contrary to how some see legal rules, natural language norms are not frustrated commands that would self-enforce if only they could find a way. If that were the goal, then wholesale formalisation and automation would make perfect sense, and there would be no need for natural language in law. That this has not happened despite many attempts draws attention to the fact that natural language plays a much larger role than simply articulating the terms of the rule. Legal rules are only effective in context, and for them to have any value in structuring society they must be interpreted at point of application, in light of that context.²⁹ It is also given legal effect in the knowledge that it cannot determine for itself in advance precisely how it will be understood, or what its meaning will be over time.³⁰ These are features of natural language that are constitutive of law; they are not bugs to be solved.

and *technological fixes* (Hart 2008); Laurence Diver, 'Computational Legalism and the Affordance of Delay in Law' (2021) 1 *Journal of Cross-disciplinary Research in Computational Law*.

²⁷ There is an analogy here with Hart's concepts of primary and secondary rules, where primary rules are directed at structuring behaviour and action, while secondary rules are about how validly to create them (see; the normativity embedded in the legal tech that frames the practice of rule-making and rule enforcement acts as another kind of 'secondary rule', or perhaps even a tertiary 'rule' (for a more detailed consideration of this idea see Diver, *Digisprudence* (n 16) 209–211).

²⁸ See 'Chapter 3. Foundational Concepts of Modern Law' in Diver and others (n 3).

²⁹ Neil MacCormick, *Rhetoric and the Rule of Law: A Theory of Legal Reasoning* (Oxford University Press 2005) ch 7. On the deep role that context and experience play in (legal) interpretation, see Hans-Georg Gadamer, *Truth and Method* (Joel Weinsheimer and Donald G Marshall trs, Bloomsbury 2013) 334ff; Stanley Fish, *Doing What Comes Naturally: Change, Rhetoric, and the Practice of Theory in Literary and Legal Studies* (Duke University Press 1989).

³⁰ Cf. Martin David Kelly, 'The "Always Speaking" Principle: Cracking an Enigma' (3 August 2023) <<https://papers.ssrn.com/abstract=4529392>> accessed 8 November 2023; Francis AR Bennion, *Understanding Common Law Legislation: Drafting and Interpretation* (Oxford University Press 2001) 17–20.

If text-driven law affords a specifically *legal* form of normativity, what kind of normativity might code produce or exert? This distinction — between code and legal normativity — is of fundamental importance to legal practice, because of the structural implications that flow from each.³¹

In stark contrast to natural language rules, rules represented in self-executing code will change the state of the world without the need for the presence or oversight of a human to interpret anything prior to its execution.³²

3.3.1 Mixing legal and technological normativity

Between these two contrasting ideas of rules, textual and computational, there are many different ways in which digital technologies, and particularly RaC approaches, can have an impact on law. The question for the purposes of this chapter is whether they tend towards supporting a *legal* idea of a rule, or a code-driven idea of a rule. To make some sense of this, the extent to which either is true can be placed on a spectrum, with notionally ‘pure’ legal normativity at one end, and notionally ‘pure’ technological normativity at the other:

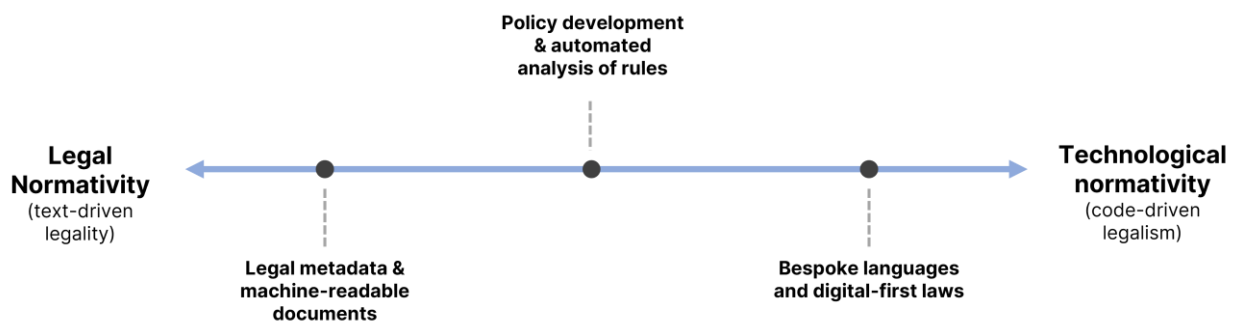


Figure 2 A spectrum of RaC normativity

Depending on the type of RaC and where it sits on the spectrum, it will to a greater or lesser degree mix the two types of normativity in one medium. Where the text of the law is embedded unchanged in a digital technology, the ‘anchoring practices’ discussed above that are afforded by text-driven law may well be preserved. But even where that is the case, the system will necessarily include *some measure* of technological normativity. It structures the behaviour and actions of those who use it, and the extent of this will depend on the RaC approach in question, and where and how it is deployed. It might exert normative force directly (e.g. via the user interface, structuring interactions with the system and framing its output³³), or indirectly, when the code-driven translations it embeds mediate the meaning of natural

³¹ For an important analysis of the distinction, see Hildebrandt, ‘Legal and Technological Normativity’ (n 22).

³² Cf. James Grimmelmann, ‘The Structure and Legal Interpretation of Computer Programs’ (2023) 1 Journal of Cross-disciplinary Research in Computational Law <<https://journalcrcl.org/crcl/article/view/19>> accessed 10 November 2023.

³³ See Andrew Le Sueur, ‘Robot Government: Automated Decision-Making and Its Implications for Parliament’ in Alexander Horne and Andrew Le Sueur (eds), *Parliament - Legislation and Accountability* (Hart 2016) 201.

language rules (e.g. where a benefits calculator provides an output that is treated as if it is legally accurate, even if this is not true).³⁴ In many cases it will do both.

Depending on whether the RaC system is used *ex ante* in the creation of legal norms, or *ex post* in the provision of advice and/or automated compliance and enforcement, the mixture of the two types of normativity will have different effects, with the technological force subsuming a lesser or greater part of what would otherwise have been legal ‘force’.

For example, the technological normativity embedded in the design of an application for drafting legislation might have some influence on the legal normativity — the natural language rules — it is designed to assist in the writing of. Similarly, the particular way the interface of a legal expert system requests answers from the citizen will frame their understanding of the system and their interactions with it. Again, the legal rules are mediated through the assumptions made by the designers of the system about what kinds of question to ask, and even how to design the question-and-answer interactions. Because those assumptions mediate the experience (one might even say ‘user experience’, or ‘citizen experience’) of the law in ways that are not neutral,³⁵ it is crucially important to consider whether or not they reflect Rule of Law values, and indeed democratic values, not least equality of access.³⁶ There can be no doubt that natural language legal texts are often obscure, complex, require expertise to understand, and can be expensive to access, and that digital technology has a role to play in solving these significant problems.³⁷ But there is at least a risk that if we attempt to tackle them by replacing the fundamentally democratic technology of natural language as the basic foundation of law, we create more long-term problems than we solve.

A central concern of COHUBICOL is what might happen when legal and technological normativity are combined in a single system, with legal force being mediated by or even converted into technological force. This is a *spectrum* of technological impact, and the picture painted above is at the extreme end where legal normativity is entirely supplanted by technological normativity. As we shall see below, that is not the goal of the vast majority of RaC approaches or their creators, although it is a vision that has been mooted by some.³⁸

Having made this brief initial foray into legal theory, which we will return to at various points throughout the chapter, we can turn to consider the spectrum of impact on law.

³⁴ For analyses of the challenges raised when code mediates the meaning and application of legal rules, see Anna Huggins, Alice Witt and Mark Burdon, ‘Digital Distortions and Interpretive Choices: A Cartographic Perspective on Encoding Regulation’ (2024) 52 *Computer Law & Security Review* 105895, 5–7; Laurence Diver, ‘Law as a User: Design, Affordance, and the Technological Mediation of Norms’ (2018) 15 *SCRIPTed* 4. For an important and damning real-world study of an example of the latter, see Rosie Mears and Sophie Howes, ‘You Reap What You Code: Universal Credit, Digitalisation and the Rule of Law’ (Child Poverty Action Group 2023) <<https://cpag.org.uk/policy-and-campaigns/report/you-reap-what-you-code>>.

³⁵ Huggins, Witt and Burdon (n 34) 7; Le Sueur (n 33).

³⁶ On those values in such contexts, see for example Monika Zalnieriute, Lyria Bennett Moses and George Williams, ‘Automating Government Decision-Making: Implications for the Rule of Law’ in Siddharth Peter De Souza and Maximilian Spohr (eds), *Technology, Innovation and Access to Justice* (Edinburgh University Press 2021).

³⁷ Cf. Lisa Burton Crawford, ‘Rules as Code and the Rule of Law’ [2023] *Public Law* 402 (discussing the deficiencies of the status quo, and the potential of digital technologies, including RaC, to ameliorate them).

³⁸ See e.g. Meng Weng Wong, ‘Rules as Code - Seven Levels of Digitisation’ (Singapore Management University Centre for Computational Law 2020) 21–23.

3.4 A spectrum of impact on law

In this section, we look at some primary classes of RaC approach, each representing a particular point of equilibrium, or disequilibrium, between legal and technological normativity. As we traverse the spectrum, the potential for more foundational impact increases:

- RaC is used to provide added information to natural language legal documents, to enhance their usefulness in terms of legal search, archiving, and knowledge management.
- RaC approaches are used as a tool to augment policy development and enhance the development of text-driven natural language rules;
- RaC translations are made available for third parties government using RaC to implement legal norms directly in their own systems in order to achieve ‘compliance’;
- The executive provides RaC translations of natural language legal rules that it uses in service delivery;
- The legislature promulgates digital-first RaC rules that are taken to be law and enforced automatically.

In considering these approaches and their potential impacts, we deliberately maintain an internal perspective on what law is and is for, as set out in depth in the first Research Study on Text-driven Law.³⁹ This means that we do not dive into the technical specifics of these RaC systems to assess how they perform on discrete computer science tasks or compared to similar approaches to the same problem. In many cases those tasks have no relationship to the goals and purposes of the law, and the notion of performance that is considered bears little relationship to legal protection and the Rule of Law.⁴⁰ Instead, therefore, we hope to highlight those claims that are particularly salient in terms of the kind of critical appraisal the COHUBICOL project aims to foster.⁴¹ This means considering questions such as: What do RaC approaches afford (and disafford), and to whom? How do they interface with or change the mode of existence of law? And what effect do they have on the capacity of the law to provide protection?

3.4.1 Stage 1: Legal metadata and machine-readable documents

We saw above that the threshold between technological normativity and legal normativity will vary according to (i) the design of a system, and (ii) the extent to which we treat its output as having legal effect. The bigger the role that technological normativity plays, the further along the spectrum the system will sit — with the potential for a simultaneous diminution in the role played by legal normativity.

At the least contentious end of the normative spectrum are RaC approaches that provide mechanisms to tag or ‘mark-up’ the structure and elements of legislative documents. Such approaches are closely connected with semantic web approaches where metadata is added to documents to provide additional information on what they contain, which in turn allows them to be processed in ways more relevant to their domain of use.

³⁹ Diver and others (n 3).

⁴⁰ Cf. P Leith, ‘Fundamental Errors in Legal Logic Programming’ (1986) 29 *The Computer Journal* 545, 100.

⁴¹ In line with the ‘method and mindset’ of the Typology of Legal Technologies (Diver and others (n 4)).

To give an example, the elements of a legislative document can be tagged to specify their structure: recitals, chapters, parts, sections, paragraphs and articles are specified as such, rather than left as blobs of text. Crucially, this structure is different from the visual structure that can be achieved in an ordinary word processor using headings, indentation and numbered lists; tagging of this kind is usually invisible to the reader.

3.4.1.1 Structural markup

One of the primary RaC technologies used for this purpose is Akoma Ntoso, an example of an eXtensible Markup Language (XML) which has been standardised and as such is freely usable by anyone — and indeed it forms the backbone of numerous legal drafting, publishing and archiving initiatives.⁴²

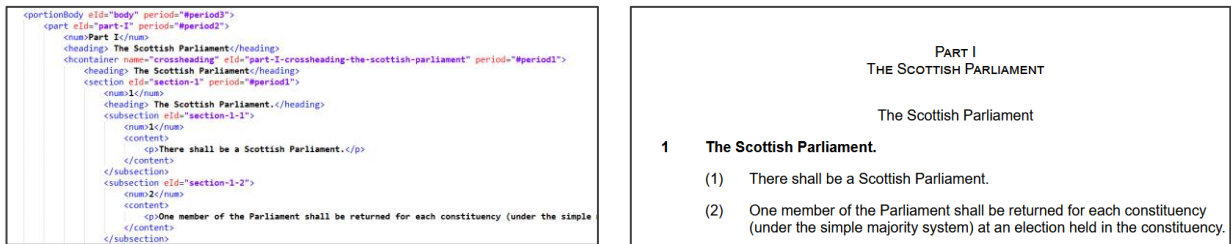


Figure 3 Akoma Ntoso markup (left) and generated HTML (right)

The left part of Figure 1 shows the markup, or tagging, within an Akoma Ntoso (AKN) version of section 1 of the Scotland Act 1998. The substantive legal text is shown in black, while the tags are shown in blue, with attributes in red. The tags specify the various elements of the Act, including its Parts, headings, sections, subsections and their numbering. From this source document further versions can be generated for human readers. The right of Figure 1 shows a web page (HTML) version of the same section, generated from the same source, with identical binding text.⁴³

At this end of the normative spectrum, the benefits of RaC are fairly unglamorous, but no less valuable for it. The fundamental structure provided by markup languages like Akoma Ntoso in turn creates a foundation which can be used to enhance other software systems that lawyers and citizens frequently rely on.⁴⁴ For example, legal search systems can utilise the structure of AKN documents to facilitate more accurate and targeted search (for example by restricting results to those provisions tagged as recitals, or pinpointing a specific individual section of an enactment). Because the document structure is explicitly specified by the creator of the tagged document, who will usually be the legislative drafter, there is less reliance on fuzzy searches that treat the content of a document essentially as a bag of words with no

⁴² See <<http://www.akomantoso.org/>> Accessed 11 November 2023. For the COHUBICOL analysis of Akoma Ntoso through the lens of the Typology of Legal Technologies, see ‘Akoma Ntoso’ in Laurence Diver and others, ‘Typology of Legal Technologies’ (COHUBICOL, 2022) <<https://publications.cohubicol.com/typology/akomantoso>>. For a short description of markup languages, see ‘Markup languages’ in ‘Computer Science Vocabulary’ (COHUBICOL, 9 October 2023) <<https://publications.cohubicol.com/vocabularies/cs/>> accessed 10 November 2023.

⁴³ Other formats can be generated from the AKN source, including PDF, RDF and generic XML.

⁴⁴ Like HTML, the markup language underpinning the Web, AKN provides a generic foundation (albeit one aimed at legal documents) upon which other systems can be built.

structure.⁴⁵ Although such search approaches have improved dramatically, the capacity to target with certainty and isolate a specific part or parts of a document and to access metadata about them is a powerfully generative affordance of this kind of RaC system. It has allowed legal archival and database systems to provide more granular access to the text of the law, and to track its evolution and status over time as it comes into force and is amended or repealed. Metadata about if and when a specific enactment has come into force allows for ‘point-in-time’ displays of which parts of a legislative instrument are in force at a given moment. The structured documents also allow for cross-references between provisions to be identified, enabling more comprehensive understanding of the interconnected legal effects of legislative provisions.

One prominent example implementing these affordances is the UK’s online public database of legislation, [legislation.gov.uk](https://www.legislation.gov.uk), which provides machine-readable versions of most primary and secondary legislation.⁴⁶ The provision of a combination of (i) authoritative legal text and (ii) a structured machine-readable format means that third parties can build applications with granular access to the official source of legal text via API.⁴⁷

3.4.1.2 Drafting software

One example of this within the public sector is the specialist legislative drafting systems that have been developed over the past decade or so. Two prominent examples of these ‘integrated development environments for legislation’ include Legislation Editing Open Software (LEOS),⁴⁸ developed by the European Commission, and LawMaker,⁴⁹ developed by the National Archives and used in the UK. Precise functionality varies, but these systems have replaced the use of word processors for drafting legislation, which required complex and unreliable template files and had limited capacity for integration with other

⁴⁵ Though there are various tools intended to parse unstructured legislative text into structured formats such as AKN. See for example Francesco Sovrano, Monica Palmirani and Fabio Vitali, ‘Deep Learning Based Multi-Label Text Classification of UNGA Resolutions’, *Proceedings of the 13th International Conference on Theory and Practice of Electronic Governance* (ACM 2020) <<https://dl.acm.org/doi/10.1145/3428502.3428604>> accessed 25 October 2023.

⁴⁶ See <<https://www.legislation.gov.uk/>> Accessed 11 November 2023. Examples of other systems that provide programmatic access to structured legislative documents include the European Union’s EUR-lex and the Australian Legal Information Institute (AustLII). Another interesting experiment currently under development by Hamish Fraser builds on the availability of structured documents to improve citation and extraction of legislative provisions. See Hamish Fraser, ‘A Love Letter to the Parliamentary Counsel of the World.’ (3 February 2023) <<https://hamish.deva-love-letter-to-the-parliamentary-counsel-of-the-world>> accessed 25 October 2023.

⁴⁷ Note that [legislation.gov.uk](https://www.legislation.gov.uk) is itself an application built on top of the underlying affordances of the computational formats, like AKN, that it utilises. This highlights the infrastructural level at which RaC technologies such as Akoma Ntoso are situated.

⁴⁸ See ‘LEOS - Open Source Software for Editing Legislation’ <<https://joinup.ec.europa.eu/collection/justice-law-and-security/solution/leos-open-source-software-editing-legislation>> accessed 11 November 2023. For the COHUBICOL analysis of LEOS through the lens of the Typology of Legal Technologies, see ‘LEOS’ in Diver and others (n 4).

⁴⁹ See ‘Legislative Drafting, Amending and Publishing Tools’ <<https://www.legislation.gov.uk/projects/drafting-tool>> accessed 11 November 2023.

systems or for collaborative working.⁵⁰ At its core, legislative drafting software produces documents directly within the structured format such as Akoma Ntoso (instead of in a generic word processor or Web format). As new elements of piece of legislation are drafted — articles, sections, subsections, paragraphs etc. — they are immediately tagged with the relevant structural elements. This is usually done ‘in the background’, such that the drafter sees only the formatted natural language document.

As more legislation is produced directly in a structured format using specialist drafting software, the requirement to convert a traditional unstructured document is removed, which in turn minimises the likelihood of errors being introduced during conversion. On top of this core difference with traditional processes of creating legislation, such drafting systems provide various additional affordances that facilitate the specialist work of legislative drafters. This includes, for example defining document structure to limit the potential for mistakes, facilitating collaborative editing across teams, tracking document versions, interlinking with legal databases for inserting/checking cross-references, sharing of drafts, and modularising common elements of the legislative workflow.⁵¹

3.4.1.3 Potential impact(s) of code-driven law

One can appreciate that by including additional metadata within the machine-readable structure of a document, approaches such as Akoma Ntoso can afford a deeper understanding of a particular piece of law which in turn might have a bearing on its legal interpretation. At the same time, there is in principle no effect on the natural language legislative document, as was shown on the right of Figure 2 — all things being equal, access to the law is not affected, nor are the traditional appearance or visual structure of the legislation. Machine-readable structured legal documents have the same interpretative affordances as do word processor or PDF versions of the same document.

3.4.1.3.1 Representing legal meaning and structure

The normative impact on law-as-we-know-it of providing basic machine-readable tagging of structure within legal documents appears minimal, provided what is tagged does not purport to provide *ex post* the *legal* meaning of the elements of the document, but rather (and only) its unambiguous structure and the metadata required to capture it. This is fundamental, because to attempt to schematise the meaning of legal norms in an unambiguous and universally-accepted way is to elide one of their core affordances: the capacity to disagree about what they ought to mean. Attempting to codify the substantive meaning of legal norms is to do law’s job before it gets the chance.

This concern is less acute with respect to the structure of legal norms or, even less problematically, the structure of the documents that contain their text. It is rarer for parties to disagree about whether a piece of text qualifies as a ‘section’ or ‘article’ of an enactment, than it is for them to argue about what that section or article ought to mean — the latter type of argument is of course the bread and butter of litigation.⁵² But potentially problematic is the inclusion of metadata about, for example, the moment at

⁵⁰ For an analysis of the affordances that an Integrated Legislative Drafting Environment, or ‘ILDE’, should have, see Elhanan Schwartz, Ittai Bar-Siman-Tov and Roy Gelbard, ‘Design Principles for Integrated Legislation Drafting Environment’ (SSRN, 30 August 2023) <<https://papers.ssrn.com/abstract=4556959>> accessed 2 November 2023. See also Diver, *Digisprudence* (n 16) 234–236.

⁵¹ These features are selected from the profile of LEOS in the COHUBICOL Typology of Legal Technologies (Diver and others (n 4)).

⁵² A case law search turned up no results for the former type of disagreement.

which a provision came into force.⁵³ If one seeks to rely on point-in-time snapshots of the state of the law, these need to be accurate to avoid potentially significant legal consequences, but dates and computers are famously not always good bedfellows,⁵⁴ and even within text-driven law the question of how to specify them reliably is not without complexity.⁵⁵ The consequences of inaccuracy could be significant, and a novel regime of liability might be necessary to account for them.⁵⁶

If it is possible for such technical issues to be overcome, and tagging of the text is limited to structural elements of legislation that are already recognised by the law, then the fundamentals of legal effect would seem to be unaffected. The mode of existence of legal norms is unchanged; legal rules are posited in natural language, they are produced by performative speech acts whose validity is governed by positive law (itself the product of the same essential process), and they become institutional facts within the legal-institutional order. The medium by which the legal texts are made available simply augments those texts with further information that may be contextually and legally relevant to their application in the real world, without affecting the capacity and methods of interpreting them. The capacity of the law to afford protection, built on the foundational capacity of text to afford multi-interpretation — and thus contest, stability, and geographical and temporal reach — are unchanged.

3.4.1.3.2 Extending the scope of legal protection

When systems are built on top of this foundation of structured documents, there is greater potential for normative impact on law and on legal effect. We saw above how searching and archival practices are extended by this type of RaC system, and how the affordances of this kind of RaC create new possibilities for legislative drafting practice. Improving the capacity to search, categorise and reference legal materials ought in principle to benefit users of the law. Affording access to the ‘raw material’ of law is a fundamental prerequisite of the capacity of citizens and their legal agents to develop novel arguments that are legally valid, a capacity that underpins legal protection.⁵⁷ This is particularly true under the conditions of contemporary law, the volume and complexity of which make it difficult if not impossible for the citizen (let alone practitioner) to make sense of all the rules applicable to a given situation. The

⁵³ This is one affordance of Akoma Ntoso, see its OASIS standard specification at ‘Akoma Ntoso Version 1.0. Part 1: XML Vocabulary’ <http://docs.oasis-open.org/legaldocml/akn-core/v1.0/akn-core-v1.0-part1-vocabulary.html#_Toc395114133> accessed 28 October 2023.

⁵⁴ On the challenges of robustly calculating dates, see for example Ana de Almeida Borges and others, ‘FV Time: A Formally Verified Coq Library’ (arXiv, 28 September 2022) <<http://arxiv.org/abs/2209.14227>> accessed 28 October 2023; Matthew Waddington, ‘Machine-Consumable Legislation: A Legislative Drafter’s Perspective – Human v Artificial Intelligence’ [2019] *The Loophole - Journal of Commonwealth Association of Legislative Counsel* 21, 46–47.

⁵⁵ See for example the guidance on dates provided by the Scottish Government to legislative drafters: Parliamentary Counsel Office, ‘Drafting Matters!’ (Scottish Government 2018) 9–12.

⁵⁶ Louis de Koker, ‘Rules as Code: The Need for an Impact Assessment to Inform Application’, *Comments on Cracking The Code: Rulemaking For Humans And Machines (August 2020 draft)* (LawTech La Trobe Research Group 2020) <<https://www.latrobe.edu.au/law/research/la-trobe-law-tech>> accessed 23 February 2022. The consequences of poor implementation can be severe, particularly for vulnerable constituencies who are perhaps more likely to be exposed to RaC systems because of the cost savings they promise. See for example the descriptions of mistranslated rules in the UK’s digitalised system for administering Universal Credit in Mears and Howes (n 34).

⁵⁷ We have argued elsewhere, particularly in the context of data-driven law, that such access to unmediated ‘lossless law’ is essential to the proper operation of the Rule of Law. See Laurence Diver and Pauline McBride, ‘Argument by Numbers: The Normative Impact of Statistical Legal Tech’ (2022) 3 *Communitas* 6.

types of computational assistance afforded by this kind of RaC system may be not just desirable, but may be *necessary*, if the Rule of Law is to have genuine purchase.⁵⁸

In that respect, then, structured legal documents afford more than unstructured legal texts and the older methods of legal search built around them. Structured documents extend what is already possible with natural language legal norms and facilitate the creation of systems that can increase and improve access to legal materials — though much will depend on the assumptions made in the design of those systems, in terms of how concepts such as relevance are handled. In principle, then, technologies at this point on the normative spectrum are an opportunity to strengthen the capacity of the law to provide protection, by facilitating more creative, forceful, or precise argumentation by reference to a wider range of relevant legal materials and the metadata that pertains to their individual status (such as validity and enforcement) and the relationships between provisions.

3.4.2 Stage 2: policy development and automated analysis of rules

The next stage on the spectrum of RaC normativity moves beyond the marking up of elements in the document to formalise additional information about the logic of those elements. The central motivating idea is that at a certain level legal rules can be abstracted *ex ante* into syllogisms: logical ‘if this, then that’ statements where conclusions flow deductively from certain premisses. The goal is then to represent the essential logic of the piece of legislation.⁵⁹

At the point of drafting legislation *ex ante*, logical representation can complement the structural markup at Stage 1 of the normative spectrum. At this stage, the goal is for computation to begin to interact with the meaning of the rules, or at least with how their meaning is likely to be interpreted once they are given legal effect.⁶⁰ At Stage 1, the tagging allows for computational tools to manipulate the metadata to provide novel affordances (better search, linked cross-referencing, granular citation, etc.). This is sometimes referred to as the rules being ‘machine-readable’.

At Stage 2, what is manipulated are symbolic representations of the rules and the relationships between them. This has been referred to as ‘machine-consumable’ rules, to highlight the different level of computational tractability. Here RaC approaches seek to capture the relationships between the symbolic representations of rules to enable conclusions to be drawn from them.⁶¹ The computer does not understand the linguistic or ‘semantic’ meaning of the rules, or their import within the legal domain — only the formal relationships between their symbolic representations. From both a legal theory and socio-

⁵⁸ Cf. Crawford (n 37).

⁵⁹ See generally Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age* (Cambridge University Press 2017) ch 2.

⁶⁰ Later we discuss what happens when computation potentially impacts the meaning of the rules.

⁶¹ For an approach that adopts defeasible reasoning to model legislative rules in the taxation domain, see Sarah B Lawsky, ‘A Logic for Statutes’ (2017) 21 *Florida Tax Review* 60. This rule-based understanding of statutory reasoning allows interim conclusions that are deduced from the rules to be ‘defeated’ by later rules, thus allowing for exceptions to be modelled. It has been influential in recent RaC initiatives, perhaps most notably in the design of the Catala language (see below).

legal perspective this is a fundamental limitation, but it does not mean such systems do not have value, for example in the production of higher-quality legislation.⁶²

3.4.2.1 Logic checking

On the basic level, logic checking can be an unglamorous addition to the drafter's toolkit, where computation augments the drafting and subsequent comprehension of the statute that already takes place. As Waddington puts it,

It can involve merely highlighting the logical structures that the drafter is trying to create in the legislation, so that any use of that logic should always be traceable, explainable and open to correction or appeal in the same way as it is when a human follows the logic from the text.⁶³

The end product is in the same medium as traditional legislative drafting — namely natural language text-driven rules — but these are the result of a process that has used some measure of testing and checking to ensure that they meet a base level of logical coherence. The result is legislation that contains fewer mistakes, for example syntactic ambiguities that create outcomes that are impossible to arrive at within the logic of the text, or cross-references to non-existent provisions.⁶⁴ Since the logic of the output is non-formal and embodied in natural language text, it can still itself be contested, along with the meaning and ascription of the predicates it contains.⁶⁵

When RaC is used in this way, the quality of the resulting legislation is higher, because when it comes to be interpreted, the chances are reduced of encountering a condition that cannot be made sense of in the real world, without recourse to a court. Although some forms of linguistic ambiguity are inherent features of natural language, and are fundamental to the contestability that lies at the heart of the Rule Law,⁶⁶

⁶² Cf. Clement Guitton and others, 'Pervasive Computational Law' (2023) 22 IEEE Pervasive Computing 48. Decades of research into symbolic reasoning in law are testament to the desire to interface computation with law in a useful way. It is outwith the scope and aims of this Research Study to canvass the full extent of this significant body of research, not to mention the philosophical assumptions on which it lies. For a useful survey see Trevor Bench-Capon and others, 'A History of AI and Law in 50 Papers: 25 Years of the International Conference on AI and Law' (2012) 20 Artificial Intelligence and Law 215. For a perspective that deeply interrogates this history from the perspective of the Rule of Law, see Gianmarco Gori, 'Law, Rules, Machines: "Artificial Legal Intelligence" and the "Artificial Reason and Judgment of the Law"' (PhD thesis, University of Florence 2021) <<https://flore.unifi.it/handle/2158/1248529>> accessed 27 October 2023. Regarding the background assumptions that have informed legal informatics research over the decades, see for example Thomas F Gordon, Guido Governatori and Antonino Rotolo, 'Rules and Norms: Requirements for Rule Interchange Languages in the Legal Domain' in Guido Governatori, John Hall and Adrian Paschke (eds), *Rule Interchange and Applications*, vol 5858 (Springer Berlin Heidelberg 2009); Layman E Allen, 'Symbolic Logic: A Razor-Edged Tool for Drafting and Interpreting Legal Documents' (1957) 66 *The Yale Law Journal* 833.

⁶³ Waddington (n 8) 182.

⁶⁴ Ashley (n 59) 46–47. Using computation for this purpose has a long pedigree in legal informatics. See for example Meldman's use of Petri nets, a kind of verifiable visual graph of a system, to model US federal civil procedure, leading to logical anomalies coming 'right to the surface': Jeffrey A Meldman, 'A Petri-Net Representation of Civil Procedure' (1977) 19 *Idea* 123, 145.

⁶⁵ Waddington (n 54) 28–29.

⁶⁶ Mireille Hildebrandt, 'A Philosophy of Technology for Computational Law' (David Mangan, Catherine Easton and Daithí Mac Mac Síthigh eds, OUP, forthcoming) <<https://osf.io/preprints/lawarxiv/7eykj/>> accessed 15 March 2021.

avoiding the drafting of patently irresolvable conditions in a statute can only be beneficial in terms of the Rule of Law, and democratically in terms of reflecting the intention of the legislator.⁶⁷

While the courtroom is where law's capacity to resolve such difficulties is most clearly demonstrated, litigation that arises from mistakes in legislation hardly reflects the higher aspirations of the law. Avoiding them in the first place frees up limited court capacity to focus on conflicts that are of more substantive importance to the individuals affected and to the community as a whole. This is true not just in terms of the closure provided by a specific ruling itself, but also in the deeper sense that the ongoing *practice* of legality is upheld, along with democratic engagement of citizens with the norms and processes that structure and co-constitute society.⁶⁸

3.4.2.2 Policy development and parallel drafting

Further along this stage of the normative spectrum, RaC can have a more structuring impact on the practices involved in developing and producing legal rules. As we saw with the definitions in section 3.2, the policy sphere is engaging with RaC as both a tool and a perspective at the interface between policymaking and legislative drafting. The OECD's working paper 'Cracking the Code' refers to RaC as "a fundamental transformation of the rulemaking process" and a "strategic and deliberate approach to rulemaking, as well as an output".⁶⁹ Waddington articulates the vision of this approach:

This could mean that legislative drafters and policy officers understand each other better during the drafting, that consultees can more easily grasp what is proposed and demonstrate how it could be changed, that inconsistencies in drafts can be spotted before they become problems, and that those who need to read the legislation can be helped to navigate complex sets of cross-references, conditions and exceptions to other exceptions. *Those would represent significant benefits in themselves, without going anywhere near automating the implementation or enforcement of legislation.*⁷⁰

The New Zealand Government's Better Rules for Government project seeks to bridge the perceived gap between policy intent and implementation, applying a service design approach to integrate policy development with rule drafting.⁷¹ It brings together teams from across the legislative process, including policy makers and analysts, legislative drafters, rule analysts, service designers and developers. Instead of following a sequential process from policy to drafting to implementation, akin to the waterfall approach in software development, the idea is that direct and iterative ('agile') collaboration between

⁶⁷ Various of Fuller's principles of formal legality in rule 'design' speak to this: laws must be reasonably clear, they ought not contradict one another, they must not require the impossible, and there should be 'congruence' between the declared rule and the state's application of it. See Fuller (n 16) ch 2. On the relationship with legislative intent, see Bennion (n 30) ch 3.

⁶⁸ Cf. e.g. Roger Brownsword, 'Lost in Translation: Legality, Regulatory Margins, and Technological Management' (2011) 26 Berkeley Technology Law Journal 1321; John Gardner, 'The Mark of Responsibility' (2003) 23 Oxford Journal of Legal Studies 15.

⁶⁹ Mohun and Roberts (n 5) 16–17 (emphasis in the original). The forms of RaC described in the previous section could be described as just 'output'.

⁷⁰ Waddington (n 8) 182 (our emphasis).

⁷¹ 'Better Rules for Government - Discovery Report' (New Zealand Government 2018) <<https://www.digital.govt.nz/dmsdocument/95-better-rules-for-government-discovery-report>> accessed 6 October 2020.

each discipline will result in rules that better reflect the policy intent of the government. As the findings of the project's Discovery Report put it:

We concluded that the initial impact of policy intent can be delivered faster, and the ability to respond to change is greater, with:

- Multidisciplinary teams that use open standards and frameworks, share and make openly available 'living' knowledge assets, and work early and meaningfully with impacted people.
- The output is machine and human consumable rules that are consistent, traceable, have equivalent reliance and are easy to manage.
- Early drafts of machine consumable rules can be used to do scenario and user testing for meaningful and early engagement with Ministers and impacted people or systems.
- Use of machine consumable rules by automated systems can provide feedback into the policy development system for continuous improvement.⁷²

Other benefits include the development of standard patterns of language to address commonly recurring policy requirements in legislation.⁷³ These set out skeleton formulations for policy aims, and the questions that must be addressed to implement the pattern in a draft. Such 'modular' rule formulations can then be used with such projects to develop software design patterns that can implement them, thus reducing the potential for misrepresentation of the rules after-the-fact, and the engineering problems of continually reimplementing what could otherwise be robust standardised approach.⁷⁴

Once again the product is still text-driven legislation, but the process by which that product is arrived at is improved in various ways, owing to closer collaboration and understanding between the cross-disciplinary teams that are involved in the policy-to-legislation process. Although the legislature cannot decide in advance the meaning of the rules it produces, their quality as *rules* is likely to be improved where the teams involved in producing them have an understanding of one another's practices and the constraints they work within, so that the 'gearbox' between democratic policymaking and legal drafting runs more smoothly.⁷⁵

⁷² *ibid* 13.

⁷³ For a set of examples and a list of legislation where they have been implemented, see Parliamentary Counsel Office, 'Guidance on Instructing Counsel: Common Legislative Solutions' (Scottish Government 2018) <<http://www.gov.scot/publications/guidance-instructing-counsel-common-legislative-solutions/>> accessed 26 May 2023.

⁷⁴ 'Better Rules for Government - Discovery Report' (n 71) 32 (suggesting RaC provides an opportunity to 'identify legislative barriers to data and digital transformation to inform the development of standard clauses, drafting guidance materials, and potential future amendments (working with PCO, DPMC Policy Project, and Stats NZ)').

⁷⁵ On the question of deciding meaning in advance, see section 3.5.3 below. For a recent discussion that seek to develop a link between legislative with software engineering processes, see Gordon Guthrie, 'Can Parliamentary and Digital Delivery Engines Ever Drive in Unison?' (*Apolitical*, 6 September 2023) <<https://apolitical.co/solution-articles/en/can-parliamentary-and-digital-delivery-engines-ever-drive-in-unison>> accessed 4 November 2023.

3.4.2.3 Potential impact(s) of code-driven law

The potential impact at this point on the spectrum depends to a large extent on whether the application of the approach is used *ex ante* during the drafting of the rules, or *ex post* in the attempt to deliver a service, automate compliance, or provide advice about what the rules mean.

As with the structural markup discussed at Stage 1, what comes out at the end of the process are rules that are, from a *formal* perspective, identical to what went before. Their fundamental affordances are unchanged (whatever their content might be). The text is still natural language, affording interpretation and contestation. The procedures of the Rule of Law are in principle still available; they take up where the legislative process leaves off, ready to deal with disagreements about right and duty in the normal way. And, where the goals of projects like the New Zealand government's Better Rules are in fact realised, the quality of the resulting rules is improved — they contain fewer logical anomalies, and their structural translation is more faithful to the goals of the legislature's policy than it might otherwise have been.

3.4.2.3.1 The threshold for formalisation

Various issues arise, however, in relation to the content of the rules. First is the question of what elements of the legislation should be logically modelled for correctness. As the authors of the Better Rules report recognise, 'not all rules are suitable for machine consumption'.⁷⁶ This raises the question of which rules are suited to computational representation and, crucially, who gets to decide this. Just as the distinction between 'easy' and 'hard' cases is not a simple one in text-driven law, the question of what rules are readily formalizable is similarly vexed. Any rule *can* be formalised; the question is *how* it is done, what the effects are of this, how they interplay with Rule of Law values and procedures, and who is affected by the change.

One of the experiments undertaken by the Better Rules project, for example, considered only one part of one piece of legislation (the rules on holiday entitlement within the Holidays Act 2003), in order to 'keep the problem (reasonably) discrete'.⁷⁷ While this is an understandable decision in an exploratory scoping exercise, the question of scalability is absolutely key, because the meaning of legal norms is never limited to just the text of the statute containing them but is influenced by other sources of law, including constitutional enactments, case law, doctrine and principle.⁷⁸

This underlines the importance of understanding the nature of text-driven law from the outset, lest the use of artificially restricted examples give the impression that success on a small scale will be applicable to the wider law and legal system. As Leith suggests, in the *legal* world (to say nothing of any other domain) it is no defence to say that the intention was to formalise only one area or piece of law, because by nature the discipline of law requires more than that:

⁷⁶ 'Better Rules for Government - Discovery Report' (n 71) 3.

⁷⁷ *ibid* 16.

⁷⁸ This incorrect presumption that legal rules can be clear and self-contained was exhibited in one of the foundational papers in legal formalisation, Marek J Sergot and others, 'The British Nationality Act as a Logic Program' (1986) 29 Communications of the ACM 370 ('the British Nationality Act is relatively self-contained, and free, for the most part, of many complicating factors that make the problem of simulating legal reasoning so much more difficult. Furthermore, at the time of our original implementation (summer 1983) the act was free of the complicating influence of case law'). See Leith, 'Fundamental Errors in Legal Logic Programming' (n 40).

legal knowledge, the sociology of law demonstrates, cannot be partitioned off into neat blocks which will fall, one by one, to the technicalism of the AI researcher. Rather, it is only by having a global appreciation of all the aspects of law which will allow each of those aspects to be properly understood — for law is an interconnected body of practices, ideology, social attitudes and legal texts, the latter being in many ways the least important.⁷⁹

Whether this limited view is a problem will depend on the context in which the RaC system is used. Policy and drafting experts using RaC in the development stage of a statute will likely understand the proper (limited) role that the system can play within their broader practice (in any event, RaC in those contexts is mostly aimed at producing well-formed documents rather than pronouncing on the meaning of the law). RaC systems that purport to facilitate compliance or give advice, however, might prompt people to take actions that are misinformed as to the proper extent and meaning of the law. This question of interpretation is of central importance, and we will return to it in section 3.5 below.

3.4.3 Stage 3: bespoke languages and digital-first laws

At this part of the normative spectrum, new domain-specific [programming] languages (DSLs) enable the declaration of rules in a format directly susceptible to computation and automated enforcement.⁸⁰ DSLs can be distinguished from general-purpose programming languages such as Python or Rust because they are designed for a particular class of computational problem or task.⁸¹

In the contemporary RaC context, the subset of DSLs known as ‘controlled natural languages’ (CNLs) are commonly adopted.⁸² These make programming rules more accessible for non-technical domain experts, i.e. policymakers and lawyers. In some cases, the language design brings the ‘grammar’ and keywords of the CNL closer to that of natural language legal text, to allow it to be read and understood as one. Despite this, as their name suggests their syntax is tightly constrained so that the rules follow a strictly predefined form.

The compilers of RaC DSLs embed different approaches to the logical aspect of legal reasoning, for example allowing for exceptions to rules (and exceptions to those exceptions), prioritisation of the applicable order of rules, and the role of time in the applicability of a rule.⁸³ The ultimate goal is basically the same: to model the logical structure of rules to produce automated conclusions that can be

⁷⁹ Leith, ‘The Application of AI to Law’ (n 16) 44. See also Frank Pasquale, ‘A Rule of Persons, Not Machines: The Limits of Legal Automation’ (2019) 87 *George Washington Law Review* 1.

⁸⁰ Cf. ‘Level five’ in Wong (n 38) 20.

⁸¹ Arie Van Deursen, Paul Klint and Joost Visser, ‘Domain-Specific Languages: An Annotated Bibliography’ (2000) 35 *ACM SIGPLAN Notices* 26 (A DSL is ‘a programming language or executable specification language that offers, through appropriate notations and abstractions, expressive power focused on, and usually restricted to, a particular problem domain.’)

⁸² Examples include *RegelSprak* and *Catala*, discussed below, and *Logical English* (Robert Kowalski and others, ‘Logical English for Law and Education’ in David S Warren and others (eds), *Prolog: The Next 50 Years* (Springer Nature Switzerland 2023) <https://doi.org/10.1007/978-3-031-35254-6_24> accessed 2 November 2023).

⁸³ See, respectively, Lawskey (n 61) (on default logic as a representation of statutory reasoning), Gerhard Brewka and Thomas Eiter, ‘Prioritizing Default Logic’ in Steffen Hölldobler (ed), *Intellectics and Computational Logic* (Springer Netherlands 2000) (on prioritized default logic), and Guido Governatori and Antonino Rotolo, ‘Changing Legal Systems: Legal Abrogations and Annulments in Defeasible Logic’ (2010) 18 *Logic Journal of the IGPL* 157 (modelling the state of rule applicability over time).

used for compliance checking, the application of the law by officials, and to provide advice on how the law might apply in a given situation (or, as is more likely, some mix of all three).

3.4.3.1 *RegelSpraak: a controlled natural language*

A contemporary example of a RaC CNL is RegelSpraak, used for the calculation of tax liabilities by the Dutch Tax Authority.⁸⁴ RegelSpraak is based on RuleSpeak, a ‘set of guidelines for expressing business rules in concise, business-friendly fashion’.⁸⁵ Like logic programming more generally, the modelling of business rules has a long history that gives an insight into the lens through which legal rules are viewed when approached from that perspective.⁸⁶ RegelSpraak imposes a strict ‘[RESULT] IF [CONDITION]’ structure on rules.⁸⁷ Those conditions compare attributes that can be Boolean (true/false), numerical, date, enumerative, or that define the role/object the rule is concerned with.⁸⁸ As with other formalisms, the rules are defined ‘atomically’ as discrete units, allowing for the identification of broader rule patterns. This is intended to mirror modularisation in programming, and allows for recursive inductive reasoning about the rules.⁸⁹

Rule pattern

Rule **<description>**
Valid from **<year>** to **<year>**
The **<attribute>** of/on a **<role | object>** must be calculated as the **<expression>**
if the **<object>** meets **<all | one | at least n | at most n | none | exactly n>** of the following conditions:
– **<condition>**
– **<condition>**

Rule

Rule **tax based on travel time**
Valid from **2011** to **2015**
The **tax** on a **train journey** must be calculated as the **TAX_PERCENTAGE** of (900 minus the **travel time by train in minutes of the taxed journey**)
if the **train journey** meets **all** of the following conditions:
– **the travel time by train in minutes is larger than or equal to 600**
– **the travel time by train in minutes is smaller than 900**

Figure 4. A RegelSpraak rule pattern (left) and substantive rule (right)

⁸⁴ Mischa Corsius and others, ‘RegelSpraak: A CNL for Executable Tax Rules Specification’, *Proceedings of the Seventh International Workshop on Controlled Natural Language (CNL 2020/21)* (2021). For our analysis of RegelSpraak through the lens of the Typology of Legal Technologies, see ‘RegelSpraak’ in Diver and others (n 4).

⁸⁵ See ‘RuleSpeak® || Let the Business People Speak Rules!’ <<https://www.rulespeak.com/en/>> accessed 2 November 2023.

⁸⁶ There is a significant literature concerned with formalising compliance between ‘internal’ business processes and ‘external’ sources such as contracts and laws. See for example Guido Governatori and Shazia Sadiq, ‘The Journey to Business Process Compliance’ in Jorge Cardoso and Wil Van der Aalst (eds), *Handbook of Research on Business Process Modeling* (IGI Global 2009). See also Guido Governatori, ‘Comments on Cracking the Code’ in ‘Comments on Cracking The Code: Rulemaking For Humans And Machines (August 2020 Draft)’ (LawTech La Trobe Research Group 2020) <<https://www.latrobe.edu.au/law/research/la-trobe-law-tech>> accessed 23 February 2022.

⁸⁷ Corsius and others (n 84) 3.

⁸⁸ *ibid.*

⁸⁹ Ilona Wilmont and others, ‘A Quality Evaluation Framework for a CNL for Agile Law Execution’, *Proceedings of the Seventh International Workshop on Controlled Natural Language (CNL 2020/21)* (2021) 6.

The aims of RegelSprak are to be intelligible to non-technical users, to allow ‘automated semantic analysis’, and to facilitate ‘automated execution of the rules’.⁹⁰ The Dutch Taxation Authority use it to automate the execution of tax rules in its internal systems and on its website, to make it easier to handle the implementation of annual budget updates, and to provide a single, centralised ‘source of truth’ for fiscal rules that, despite their quasi-natural language representation, are “only interpretable in one way”.⁹¹ Given its intended audience, its implementation within the Dutch Taxation Authority (DTA) uses Dutch phrasing and grammatical construction to ‘maximize the resemblance to a natural sentence’.⁹²

3.4.3.2 Catala: a domain-specific language

Another prominent example of a DSL used in RaC is Catala, a language originally designed for use in applying French fiscal law.⁹³ It adapts the ‘literate programming’ approach to code documentation by putting the executable RaC code immediately adjacent to the legal text which it seeks to translate.⁹⁴

```
104     ### (3) Employee discount defined
105
106     The term “employee discount” means the amount by which-
107
108     (A) the price at which the property or services are provided by the employer to
109     an employee for use by such employee, is less than
110
111     (B) the price at which such property or services are being offered by the
112     employer to customers.
113
114     ```catala
115     scope QualifiedEmployeeDiscount:
116         assertion customer_price >= employee_price
117
118         definition employee_discount equals
119             customer_price - employee_price
120     ```
```

Figure 5 Part of the US federal tax code formalised in Catala⁹⁵

Here, literate programming it is not (just) intended to aid understanding of the code ex post, but is also integral to the ‘pair programming’ approach used to writing that code. There, an expert in Catala sits alongside an expert in the relevant fiscal law. Again, the cross-disciplinary approach to policy development discussed above can be facilitated here: expertise in policy, programming and legal interpretation can be mutually beneficial for the end product, the legal or policy expert interpreting the

⁹⁰ *ibid* 2.

⁹¹ Frans Fokkenrood, ‘RegelSprak for Business Rules: Experiences in Building a Business Rules Compiler for the Dutch Tax Administration’ (2011) 12 Business Rules Community <<https://www.brcommunity.com/articles.php?id=b622>> accessed 6 November 2023.

⁹² Corsius and others (n 84) 2.

⁹³ Denis Merigoux, Nicolas Chataing and Jonathan Protzenko, ‘Catala: A Programming Language for the Law’ (2021) 5 Proceedings of the ACM on Programming Languages 1, 9. For our analysis of Catala through the lens of the Typology of Legal Technologies, see ‘Catala’ in Diver and others (n 4).

⁹⁴ For the classic discussion, see Donald Knuth, ‘Literate Programming’ (1984) 27 The Computer Journal 97.

⁹⁵ This is an indicative example, reproduced from Denis Merigoux and others, ‘Catala’ <<https://github.com/CatalaLang/catala>> accessed 12 November 2023.

legal norms and checking the programmer's implementation in a cycle of iterative testing and refinement.⁹⁶

The use of a DSL adds an extra layer to the RaC-enhanced policy development process discussed above at Stage 2, by providing a formalisation into which the text of the legal norms can be directly translated. The result is a text that is bi-directional; it is both a natural language text, intelligible by humans without a technical background, and a formalisation that can be compiled and executed by the machine.⁹⁷ Where the compiler of the DSL is designed around formal verification principles, as Catala is, one can be certain that the compiled code reproduces the logic of the rules expressed in the DSL.⁹⁸

This builds on the logic *checking* capacity of logic programming to facilitate the general-purpose code that can be used in production systems. The translation of the logical structure of the legislation into a medium that can be (i) checked computationally, and (ii) understood by a non-technical reader, means that the latter — usually a legal expert — can verify that the output of the code is congruent with how they expect the legal rules to operate.

3.4.3.3 A 'single source of truth'

The compiler of the DSL converts the intelligible text of the DSL model into a form that allows the modelled logic to be computed just like any other programme. Some languages, such as Catala, are designed to allow trans-compilation of that logic into general-purpose programming languages. The resulting code can then be integrated directly into user-facing applications.⁹⁹ To the extent that the application is required to comply with a specific law, or purports to enforce or give advice on what the law means, it can be said to be faithful to the model of the law that was written in the DSL.

The goal here is to create a 'single source of truth':¹⁰⁰ a quasi-natural language version of the law that is endorsed by legal experts as a canonical translation of the text-driven law that it models. This is somewhat analogous to the structured versions of legislative documents we saw at Stage 1 of the normative spectrum. Like those structured documents, the initial DSL translation is generative, in that it can be used as a source from which to produce multiple further versions for use in different contexts. The significant difference at this point on the spectrum, however, is that what can be done with those versions is potentially more impactful because they are computable, automatable, and because they can be directly integrated into infrastructural or application code that will operate *ex post* in the real world.

⁹⁶ Merigoux, Chataing and Protzenko (n 93) 22; Denis Merigoux and Liane Huttner, 'Catala: Moving Towards the Future of Legal Expert Systems' (INRIA 2020) <<https://hal.inria.fr/hal-02936606>> accessed 25 February 2021.

⁹⁷ For an analysis of what the bi-directionality means for different parties, see '3.1 Code's Bi-directionality: A Text for both Computers and Humans' in Laurence Diver, 'Interpreting the Rule(s) of Code: Performance, Performativity, and Production' [2021] MIT Computational Law Report <<https://law.mit.edu/pub/interpretingtherulesofcode/>> accessed 30 October 2023.

⁹⁸ Denis Merigoux, Raphaël Monat and Jonathan Protzenko, 'A Modern Compiler for the French Tax Code' (ACM 2021) <<https://inria.hal.science/hal-03002266>> accessed 3 November 2023. This will also be true for trans-compiled code (see below).

⁹⁹ Cf Merigoux, Chataing and Protzenko (n 93) 25, taking this approach to demonstrate a prototype for a web-based benefits calculator. See also Merigoux, Monat and Protzenko (n 98).

¹⁰⁰ Wong (n 38) 5. See also the requirement specified by the policy stakeholders for the design of RegelSprak: Corsius and others (n 84) 2.

The threshold between technological normativity and legal normativity is thus quite different from RaC approaches at Stage 1 on the spectrum.

3.4.3.4 *Digital-first laws*

In the approaches outlined above, the goal is to reduce the friction of translation as far as possible, so that those involved in developing policy and the natural language rules that implement it can also be directly involved in the process of producing RaC translations for use in digital service delivery, compliance, or legal advice and knowledge management.

Going further, the direction of travel is moving towards various visions of ‘digital-first’ drafting, including the direct use of DSLs for *legal* rules.¹⁰¹ Policymaking is oriented around simplifying rules and disambiguating terms, in order to reduce discretion and facilitate automated processing.¹⁰² Here the *digital* version is treated as the ‘single source of truth’. The law is expressed directly in — rather than translated into — an executable, presumably still human-readable, DSL, and thus compliance is essentially guaranteed, provided the rules are successfully integrated into the target systems. Friction between articulating a rule and it being implemented in a digital system is reduced as far as possible and, in an inversion of what was discussed above, natural language versions of the rules are generated *from* the DSL.

Painting what is probably an extreme picture, Wong envisages that at this point legislative drafters would consider the digital RaC version to be authoritative, while the public would interact with the natural language version that is generated from the latter. Taken further still, the digital version comes to be treated as the official source of law both inside the administration *and* by the public.¹⁰³ Both legislative and contractual norms are digital-by-default; just as digital audio, video and image formats have become the default means of representing previously ‘analogue’ media such as music, film and imagery, so too can the law have its essential substance represented digitally.¹⁰⁴ If the legislature was to get to this point, where law is represented directly in code, the layers and steps of the legislative process would

¹⁰¹ See for example the efforts of the European Commission’s ‘Better Legislation for Smoother Implementation’ (BLSI) project (‘Better Legislation for Smoother Implementation | European Commission’ (27 October 2023) <<https://joinup.ec.europa.eu/collection/better-legislation-smoother-implementation>> accessed 9 May 2023) and its associated SEMIC conference, and the Danish Government’s digital-ready legislation project (‘Digital-Ready Legislation | Agency for Digital Government’ <<https://en.digst.dk/digital-transformation/digital-ready-legislation/>> accessed 17 May 2023).

¹⁰² Ursula Plesner and Lise Justesen, ‘The Double Darkness of Digitalization: Shaping Digital-Ready Legislation to Reshape the Conditions for Public-Sector Digitalization’ (2022) 47 *Science, Technology, & Human Values* 146, 147.

¹⁰³ Wong refers to first to ‘Code = Natural’ and then ‘Code > Natural’ (n 38) 21).

¹⁰⁴ *ibid* 23. This vision draws implicitly on the abstractions that underpin information theory (cf. the seminal analysis in Claude E Shannon, ‘A Mathematical Theory of Communication’ (1948) 27 *The Bell system technical journal* 379). As will be discussed in section 3.5, it does not and cannot capture the entirety of the law, in particular its institutional nature: see Mireille Hildebrandt, ‘Law as Information in the Era of Data-Driven Agency’ (2016) 79 *The Modern Law Review* 1.

be dramatically reduced, and the code-driven rules would presumably have the imprimatur of constitutional validity.¹⁰⁵ We would have Rules as Code as Law.

3.4.3.5 Potential impact(s) of code-driven law

At this point on the normative spectrum, we are starting to depart from legal normativity *per se*. Constrained natural languages and domain-specific languages that are human-readable but machine-executable limit what can be validly expressed in them, which can have a shaping effect on the content of policy from the outset.

Even if the scope of policy is left unchanged (assuming that is possible in the shift from text to code — see the discussion at section 3.5.3 below), and all that is aimed for is a faithful or ‘isomorphic’ representation of the logic of the statute, this remains because of the way technological normativity operates as compared to text-driven normativity.¹⁰⁶ At the point of drafting natural language rules logic modelling can be helpful, as we have seen above, but at the point of execution or the provision of advice it necessarily elides large parts of what it is to interpret and apply a legal rule, and likely too the procedures that implicitly accompany text-driven normativity. For that reason, ‘isomorphism’ as a goal is fundamentally limited in ways that circumscribe its usefulness for real-world application of legal rules. Even a robust, formally-verified isomorphism must not be confused with the law itself, since rules, however high-quality they might be, are not the whole of the law, even when they are text-driven.

If rules are written directly into code the effect is stronger still. In the latter case the disconnect between legal and technological normativity is complete: even if a natural language text is derived from the code and treated as notionally ‘authoritative’,¹⁰⁷ in practice the normative divergence would mean those subject to the automated execution of the rules would be interpreting a vision of normativity (textual/legal) that is categorically different from that which would be imposed in reality (i.e. technological). This would be deeply problematic in terms of the Rule of Law; the executive and the public would not be ‘reading from the same hymn sheet’, normatively speaking, undermining the Rule of Law principle of equality, and the affordance of procedural due process. Those with access to the code would have a different view of what the law is from those with access only to the natural language version. The democratic affordances of textual interpretation would be undermined, in favour of those who have access to the code and hence prior knowledge of the distinct technological normativity it will impose.

Articulating laws directly into code for the purposes of execution shifts us closer to computational legalism. It sidesteps the affordances of text-driven law, because automation is precisely the goal, eliding the direct and indirect values of those affordances. The relationship between legal and technological

¹⁰⁵ Tom Barraclough, Hamish Fraser and Curtis Barnes, ‘Legislation as Code for New Zealand: Opportunities, Risks, and Recommendations’ (Brainbox; The Law Foundation 2021) para 254 <https://www.brainbox.institute/_files/ugd/13cbd1_cf3fd1723fb547c1ac00310ad20c0781.pdf> accessed 24 March 2022. In the UK, the Interpretation Act 1978 and Interpretation and Legislative Reform (Scotland) Act 2010 do not explicitly define legislation as textual. Multiple references are made to ‘enactments’ and ‘provisions’, but apart from various references to ‘words’, the nature of these is implied rather than specified.

¹⁰⁶ Gordon, Governatori and Rotolo define isomorphism as ‘a one-to-one correspondence between the rules in the formal model and the units of natural language text which express the rules in the original legal sources’ (n 62) 285.

¹⁰⁷ Wong (n 38) 21.

normativity shifts dramatically away from an equilibrium between rules posited *ex ante* and interpretation and procedure happening *ex post*. Rules become the central focus, and at the same time because those rules are *computational* rather than simply textual, and they impose *technological* rather than *legal* normativity, their capacity to impose themselves as much stronger, and therefore their constitutional acceptability is much weaker.¹⁰⁸

3.5 Anticipating legal protection under code-driven law

Having set out a spectrum of potential impact on law of RaC, in this section we turn to broader concerns with regard to code-driven law and legal protection. We begin with some general considerations, before turning to the fundamental issues of interpretative authority and the effect on legal effect.

3.5.1 Who benefits?

An important question that must be answered with respect to each RaC application is: who benefits? The intention generally is to improve access to justice, reduce cost, and/or more efficiently convert policy into legal rules. It will take some time to tell whether these laudable goals have been realised, but whatever the answer it remains the case that formalisation means adopting a certain view of the world. At least when used to execute the law, or advise on its content, it imposes a framing that it ought to be the law's role, in a plural society, to keep open. As Schafer puts it,

Legal AI becomes the stalking horse of a very specific conception of justice, turning what should be a contested public debate about the values of law into a technocratic decision of what is computationally possible.¹⁰⁹

For all its faults in practice — some of which can indeed be ameliorated by RaC — text-driven law is fundamentally democratic insofar as natural language and the normativity that accompanies it afford both accessibility and the co-existence of a multitude of differing worldviews. The systems and procedures built around the technology of text might be flawed and in need of (serious) reform, but that is not a consequence of text as the central medium of law.

Regarding access to justice, we suggested above that certain RaC approaches, such as structured documents, can materially enhance access to the text of the law, and provide affordances that are genuinely valuable for interpreting legal meaning and legal status. But other RaC systems threaten to create a two-tier justice system, where those without power and the means to access bespoke legal advice instead have to make do with commoditised output of a formalism. Those people will in some cases be in vulnerable positions and, given the normativity of the computational medium, might not appreciate how to manipulate the law to fit their needs or understand what their options are for

¹⁰⁸ Cf. Marco Goldoni, 'The Politics of Code as Law: Toward Input Reasons' in J Reichel and AS Lind (eds), *Freedom of Expression, the Internet and Democracy* (Brill 2015); Bert-Jaap Koops, 'Criteria for Normative Technology: The Acceptability of "Code as Law" in Light of Democratic and Constitutional Values' in Roger Brownsword and Karen Yeung (eds), *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes* (Hart 2008); Diver, *Digisprudence* (n 16) ch 3.

¹⁰⁹ Burkhard Schafer, 'Legal Tech and Computational Legal Theory' in Georg Borges and Christoph Sorge (eds), *Law and Technology in a Global Digital Society* (Springer International Publishing 2022) 320 <https://link.springer.com/10.1007/978-3-030-90513-2_15> accessed 14 June 2022.

contestation.¹¹⁰ The risk is that existing problems of access to justice are potentially amplified rather than solved, with efforts better placed for example toward providing greater access to legal aid (more costly than RaC and therefore less attractive to some though this is likely to be).

Lastly is the question of where lawyers' skills really lie. A great deal of what lawyers actually do lies outside the bare interpretation of rules.¹¹¹ By focusing on rules and their automation, we potentially deskill lawyers and reduce the extent of their role as skilled interpreters of those rules in light of their clients' circumstances, and those circumstances in light of the rules.¹¹² A commodification of expertise might remove lawyerly practices that are societally valuable both in terms of providing high-quality legal advice, but also independently of the outcome of a case, for example the giving of support, understanding, and solidarity.

At the same time, it is possible that automation might instead free up time for those elements of the role. Understanding the impact will require further (empirical) research.¹¹³ What is necessary, however, is a close(r) understanding of what legal protection and the Rule of Law actually need in order to operate, and thus what practitioners require. Questions of 'user need' might in fact be *better* answered in part by the design thinking approaches adopted by projects like Better Rules.¹¹⁴ To that extent they might be welcomed as a means of getting to the heart of how those who 'do' law can be better understood and supported by technology to uphold its central values.

3.5.2 Increased complexity and maintenance

It is possible that, rather than reducing complexity, some RaC approaches will produce more of it, or even require it in order to remain feasible. There may be ripple effects in the legal system, depending on how RaC is adopted — particularly if its outputs are treated as having legal effect.¹¹⁵ For example, procedure and due process might need to be adapted to account for the speed of RaC outputs. Translations in one part of the system might necessitate a cascade of translations in other parts, if we are to avoid the complications of attempting to combine legal and technological normativity in or around the same subject matter.¹¹⁶ Areas of law that might hitherto have been thought to be inappropriate for formalisation might come to *require* translation, for example interpretative and procedural provisions, in

¹¹⁰ Cf. the experience of benefit claimants in 'You Reap What You Code' by Mears and Howes (n 34). See also the discussion of street-level bureaucracy below.

¹¹¹ John Morison and Philip Leith, *The Barrister's World and The Nature of Law* (Open University Press 1992); Philip Leith, 'The Rise and Fall of the Legal Expert System' (2016) 30 *International Review of Law, Computers & Technology* 94, 101–103.

¹¹² See Tatiana Duarte, '3.5.2 Legal Reasoning and Interpretation', *Research Study on Text-Driven Law* (COHUBICOL 2023) <<https://publications.cohubicol.com/research-studies/text-driven-law/chapter-3/legal-reasoning-and-interpretation/legal-reasoning-and-interpretation/>> accessed 18 September 2023; Philip Leith, 'The Problem with Law in Books and Law in Computers: The Oral Nature of Law' (1992) 6 *Artificial Intelligence Review* 227.

¹¹³ Leith suggests that empirical research consistently finds that in practice the legal expert systems of the last generation did not provide lawyers and judges with much assistance. See Leith, 'The Rise and Fall of the Legal Expert System' (n 111) 101 ('their needs are not met by a system which simply lists rules and indicates the ordering in which they were triggered.')

¹¹⁴ For an exploration of design thinking in the legal domain, see Rae Morgan, 'Lawyers Are Still Lawyers. Except When They're Not.' in Emily Allbon and Amanda Perry-Kessaris (eds), *Design in Legal Education* (Routledge 2022).

¹¹⁵ See section 3.5.4 below.

¹¹⁶ The next section discusses some of the problems of mixing legal and technological normativity.

order to support those areas whose formalisation is thought to be uncontroversial. Alternatively, if they are still thought to be resistant to translation, a kind of parallel set of equivalent code-driven rules might need to be developed, just to support the code-driven parts of the law. Interfaces will be required to connect the code-driven body of law with the textual, particularly where the latter is deemed authoritative.¹¹⁷ The potential complexity of the interplay between textual rules and code-driven rules is something that will need to be properly anticipated.

Another area of potential complexity is the maintenance of the rules: will they be kept up to date, and by whom? While we generally accept that officially published legislation is often not kept perfectly current, as we have repeatedly seen technological normativity is quite different in its capacity to impose itself, and so the problem of inaccurate or out of date rules becomes hugely salient. RaC interpretations that do not adequately reflect the state of the law in light of amendments or repeals, or do not reflect the interpretations of that law by the courts or in light of other instruments that have a bearing on their meaning (see the last paragraph), risk producing outputs that are readily relied upon because of the medium that delivers them but which are legally invalid. If text-driven norms are ‘always speaking’, even if purposive interpretation allows for adaptation to meet new or unforeseen circumstances, the potential risk with inaccurate RaC translations is that they are ‘never listening’.¹¹⁸ Reliance on such translations, which might be inadvertent if they are automatically imposed, could have significant consequences.

3.5.3 Interpretative authority

Decisions about (i) what gets formalised, and in what ways, (ii) who makes that choice and under what authority, (iii) where the resulting system will be deployed and for what purposes, and (iv) which citizens and legal subject is it aimed at, have significant normative implications.¹¹⁹

Where a deployment purports to enforce the law or provide advice as to what it means, it is not sufficient to rely on the coders of RaC translations to ensure there is a ‘human-in-the-loop’.¹²⁰ This puts the constitutional cart before the horse: not only do those who write the rules get to decide on their interpretation and how even the notionally incontestable logic of the rules should apply, they also decide on the ‘escape hatches’ for when something goes wrong.

This places too much power in the hands of those creating the rules, undermining the separation of powers and the role of the court in providing authoritative interpretations of the meaning of rules in particular cases (in the knowledge and foresight that those interpretations will have salience in future analogous cases). As Bennion puts it,

It is the function of the court along authoritatively to *declare* the legal meaning of an enactment. If anyone else, such as its drafter or the politician promoting it, purports to lay down what the legal

¹¹⁷ We saw above in section 3.4.5.1 (‘The mirage of human-readable code’) why it is problematic to speak of natural language rules as authoritative when the implementation is direct via code-driven rules, and is thus implemented via an entirely different kind of normativity.

¹¹⁸ Kelly (n 30); Bennion (n 30) ch 3.

¹¹⁹ It is notable that in its recommendations the Better Rules report refers to the overlapping interests of various agencies, including the executive’s policy and service innovation teams, the Parliamentary Counsel’s Office, and the Internal Revenue department, but does not explicitly mention citizens or the courts in that context.

¹²⁰ Wong suggests, for example, that ‘designers of systems are exhorted to leave in entry-points for human discretion’ (Wong (n 38) 23).

meaning is the court may react adversely, regarding this as an encroachment on its constitutional sphere.¹²¹

It also obscures the responsibility for ascribing or attributing a particular meaning to the rule in a particular case, and thus the justificatory element that is (or ought to be) inherent in any application of a legal rule — along with which comes the responsibility to provide reasons for the conclusion that was reached. A significant chunk of the reasoning is front-loaded, in the belief it is deductively universal and therefore not part of the open texture of the law.¹²² Compliance is foretold; legal subjects are objects of control, rather than agents who get to choose to comply, and how. Engagement with the community might be stunted if we no longer must actively interpret the relevance of rules within a given context.¹²³

3.5.3.1 *The mirage of human-readable code*

As briefly mentioned above, the use of quasi-natural language runs the risk that the body of rules that is developed gets shifted to suit what can be represented in the DSL, even within the policy development process. The fact that it looks much like natural language heightens this risk; some might believe that because it looks like natural language, anything can be formalised in it, or, conversely, that anything that cannot be formalised in it is not worth including in the law. This is a framing effect that over time risks limiting the scope of substantive legal protection. It raises the question of whether RaC rules should in fact *avoid* being human readable.

In the *ex post* context of application, given that the automated execution of digitised rules is of a nature categorically different from how a textual rule is ‘executed’, making RaC rules look as close to natural language rules as possible might in fact mislead as to their nature. This is not an argument against the requirement for transparency or explanations/evidence about how they have operated in practice; affording these can be achieved for example through the better design of automated logging. The risk of blurring the line between legal normativity and technological normativity might imply that in fact digital representations of the law, insofar as they are directed at *ex post* application by and to citizens, should actively seek to *avoid* appearing too similar to natural language laws.

This ambiguity is deepened when we consider how judges should respond to gaps in digital-first RaC laws.¹²⁴ How would they fulfil their constitutional function in a text-driven way, when the original norm is code-driven? If judges produce a natural language judgment about a space left in a RaC translation, does that then need to be converted into additional code and added to the RaC implementation? Would the separation of powers require that the judges produce the code themselves? Or will they write orthodox (i.e. textual, legal normative) orders requiring RaC coders to amend the digital translation? In that case, which seems most plausible, we come back round to the problem of interpretative authority — the constitutionally-empowered court decides, but it is the RaC coder, probably within the executive, who interprets that judgment and implements it in their system. We might end up in an infinite regress, with no acceptable closure.

¹²¹ Bennion (n 30) 17.

¹²² Pasquale (n 79) 4–5.

¹²³ Roger Brownsword, ‘Code, Control, and Choice: Why East Is East and West Is West’ (2005) 25 *Legal Studies* 1.

¹²⁴ Cf. Felicity Bell and others, ‘AI Decision-Making and the Courts: A Guide for Judges, Tribunal Members and Court Administrators’ (The Australasian Institute of Judicial Administration Incorporated 2022) 29.

This emphasises the point that RaC systems should never be given legal effect, assuming we still wish the courts to have a text-driven adjudicative role. To do so would introduce logical contortions like the one just described. It would undermine the fabric of the law and the way that legal rules fit into, and reflexively constitute, a complex web of interrelated practices that are oriented towards legal protection and the Rule of Law.

It might be that RaC formalisms should therefore emphasise their technological character, rather than seeking to ape natural language, in order to highlight that they are tools of implementation rather than canonical sources of law. The difference between the two forms of normativity must at all times be clear.

3.5.3.2 Technological normativity and interpretation

Legislative drafters are enjoined to write laws that respect and uphold the Rule of Law; in seeking to produce an ‘internally coherent conceptual scheme’ of rules, they serve ‘core rules of law values of legal certainty, predictability, formal justice and equality.’¹²⁵ As we saw above, law cannot be split into discrete self-contained parts, and by the same token drafting requires a sensitivity to the broader legal domain within which a new piece of legislation will sit, adapting terminology and conceptual structure to ensure coherence with what has gone before.

If the drafter fails to do this, or does it badly, the problem could be solved through interpretation.¹²⁶ While legal reasoning can ultimately be presented as syllogistic logic, that comes only at the point of justification of an argument, after various interpretative hurdles have been passed and attributions made (some of which might result from arguments about what that logic itself ought to be), not before.¹²⁷ The gap between interpreting a text rule and following through on its implications means latent incoherences in the text can be identified, ignored or if necessary contested. This inherent passivity is central to the nature of text-driven law and the spaces it affords for considered action.¹²⁸

Code, however, is different. Its execution is effectively immediate, and clear-edged.¹²⁹ When used for enforcement or advice, the output is the output, echoing the legalist idea that ‘the law is the law’.¹³⁰ If the coded model fails to integrate properly with models of other legislation that are relevant to it, it will simply fail to execute as expected — perhaps without anyone being aware of it. Depending on the extent to which the output is treated as having legal effect, the consequences of this type of normativity will vary. As Barraclough, Fraser and Barnes emphasise, RaC representations ought therefore to be

¹²⁵ Philip Sales, ‘The Contribution of Legislative Drafting to the Rule of Law’ (2018) 77 *The Cambridge Law Journal* 630, 633.

¹²⁶ Bennion (n 30) chs 3–5.

¹²⁷ MacCormick (n 29) chs 3–4 (At 71: ‘The conclusion to be drawn for all cases, the legal one included, is not that ascriptive decisions or determinations preclude or exclude deductive logic, but rather that they are a necessary precursor to any deductive reasoning whatsoever that is carried out with reference to the actual world. Every form of applied logic requires decisions as to the applicability of universals (predicate terms) to particular instances... If these require justification in given pragmatic circumstances, then certainly that “external” justification has to be provided before any syllogistic representation of a conclusion can be convincing.’).

¹²⁸ Diver, ‘Computational Legalism and the Affordance of Delay in Law’ (n 26).

¹²⁹ These are central aspects of ‘computational legalism’ that make technological normativity more problematic than even text-driven legalism (*ibid.*).

¹³⁰ Cf. G Radbruch, ‘Five Minutes of Legal Philosophy (1945)’ (2006) 26 *Oxford Journal of Legal Studies* 13.

thought of not as *translations* of the law, but instead as individual *interpretations* of it.¹³¹ This highlights that any given RaC model is not the law *per se* but is just one interpretation of what the law says.¹³²

But the role played by technological normativity creates a crucial difference between code-driven interpretations and other non-authoritative (i.e. non-judicial) interpretations. People, including lawyers, interpret the law all the time, and often they will get it wrong. But the difference with code-embedded interpretations and the systems that execute them is that the ‘user experience’ of the citizen is quite different: the embeddedness of the interpretation tends away from the capacity of those affected by it to question its validity or applicability. Depending on whether and how the RaC translation’s (lack of) authority is communicated, they might be misled into accepting the output as accurate or binding. In such situations the problem lies as much in the design of the application as it does in the specific formalism that is used; as Le Sueur suggests, the application in a very real sense becomes part of the law:

we should treat ‘the app’ (the computer programs that will produce individual decisions) as ‘the law’. It is this app, not the text of legislation, that will regulate the legal relationship between citizen and state in automated decision-making. Apps should, like other forms of legislation, be brought under democratic control.¹³³

In this vein, Waddington makes a salient distinction between a RaC system answering the question ‘what does the [original] act say’ on the one hand, and ‘what does the act mean (what are my rights)’ on the other.¹³⁴ The difference is key, and whether or not it is understood by the users of a RaC system is in large part a question of design. As he observes, most jurisdictions publish legislation without later including alongside it all the relevant caselaw that has a bearing on its legal meaning. Pragmatically, if RaC outputs are treated as authoritative, then that is what they are, at least for those who take action based on their output. If a RaC system is taken to be giving concrete advice rather than discrete, non-authoritative interpretations, the system has overreached. By the same token, if the limit of the interpretations’ authority is properly expressed at the point of interaction, in order to communicate that the output does not constitute advice or have legal effect, then insofar as the system benefits citizens’ understanding of their legal status it might be considered beneficial in terms of enhancing legal protection.¹³⁵ Communicating this adequately is a question of design. Whether or not it will have a bearing on the perceived utility of the system remains to be seen.

3.5.3.3 Technological normativity and discretion

We saw above that in some cases an explicit goal of RaC is the reduction of discretion and the promotion of automation. As Plesner and Justesen describe the Danish approach,

¹³¹ Barraclough, Fraser and Barnes (n 105) pt 3 (‘Code should not be legislation’), 42-79.

¹³² Cf. the discussion on the effect on legal effect in section 3.5.4 below.

¹³³ Le Sueur (n 33) 201.

¹³⁴ Waddington (n 8) n 17.

¹³⁵ This is the approach taken by the legal expert system DataLex, for example, where the individual ‘consultations’ display prominent disclaimers that they should be used ‘only for education and testing purposes’, and the user must agree the output ‘will not be relied upon for any purpose’. See ‘DataLex: AustLII’s Legal Reasoning Application Platform’ (n 13).

whenever possible, legislation should build on simple rules and unambiguous terminology to *reduce the need for professional (human) discretion*, thereby allowing for the extended use of automated case processing across all types of public-sector organizations and policy areas.¹³⁶

Lipsky's concept of 'street-level bureaucracy'¹³⁷ is relevant here, and reflects the dignitarian idea, inherent to legality, that law ought to be applied with sensitivity to particular contexts and needs. He suggests that those who are involved in the implementation of legal rules 'on the ground' invariably use their contextually informed judgement, whether they are benefits administrators, the police or even the judiciary.¹³⁸ Discretion shapes the interactions between such actors and citizens, which are informed by much more than the bare terms of the rule and the abstract status of the citizen. It is worth quoting Lipsky at length:

The essence of street-level bureaucracies is that they require people to make decisions about other people. Street-level bureaucrats have discretion because the nature of service provision *calls for human judgment that cannot be programmed and for which machines cannot substitute*. Street-level bureaucrats have responsibility for making unique and fully appropriate responses to individual clients and their situations... the unique aspects of people and their situations will be apprehended by public service workers and translated into courses of action responsive to each case within (more or less broad) limits imposed by their agencies. They will not, in fact, dispose of every case in unique fashion. *The limitations on possible responses are often circumscribed, for example, by the prevailing statutory provisions of the law or the categories of services to which recipients can be assigned*. However, *street-level bureaucrats still have the responsibility at least to be open to the possibility that each client presents special circumstances and opportunities that may require fresh thinking and flexible action*.¹³⁹

This highlights the distinction between what bare laws require on the one hand, and the reality of their application in complex real-world contexts on the other. This resonates with the distinction between legalism and legality; the automatic application of bare rules absent reasoned interpretation on the one hand, versus choosing whether and how to apply a rule in light of context on the other.¹⁴⁰ The requirement to treat individuals with individual dignity and autonomy is a core tenet of liberal legality, and though it may not be achieved as often as it should be, seeking to restrict the discretion that is a necessary part of realising it is not a solution.

A recent report by the Child Poverty Action Group on the UK's digitised Universal Credit benefit provides a useful case in point.¹⁴¹ The authors concluded that the digitalised implementation of the benefit

¹³⁶ Plesner and Justesen (n 102) 3 (our emphasis). The pursuit of 'simple rules and unambiguous terminology' connects back to the discussion above about policy being shaped by the RaC medium.

¹³⁷ Michael Lipsky, *Street-Level Bureaucracy: Dilemmas of the Individual in Public Service* (30th Anniversary Edition, Russell Sage Foundation 2010). See also Le Sueur (n 33) 192–193.

¹³⁸ Lipsky lists as typical street-level bureaucrats 'teachers, police officers and other law enforcement personnel, social workers, judges, public lawyers and other court officers, health workers, and many other public employees who grant access to government programs and provide services within them.' (Lipsky (n 137) 3).

¹³⁹ *ibid* 161.

¹⁴⁰ Leith notes that in the context of applying welfare rights there is 'an attempt to keep away from legalism and legal rules as much as possible'. See Leith, 'The Application of AI to Law' (n 16) 43.

¹⁴¹ Mears and Howes (n 34). Leith adverted to problems with this precise application of legal expert systems, as far back as 1988: 'these tactics [computerising social security] would not simply be applying computers to ease

undermined Rule of Law principles of transparency, procedural fairness, and lawfulness. They found that the interpretation of the law embedded in the system diverged from the terms of the various pieces of underlying legislation. Perhaps more importantly, the design of the system did not provide the latitude for ‘street-level’ discretionary ‘work arounds’ that might have solved some of the problems caused by those mistranslations, for example the capacity to submit claims earlier than usual in certain circumstances where this is justified on grounds of fairness. The effect is that those who perhaps need the protection of the law the most were denied it. The authors noted that this was not an inevitable consequence of digitalisation, and with careful design choices some of the system’s pitfalls could have been avoided.¹⁴²

With the accelerating shift towards the ‘digital state’, and with it the delivery of services via RaC-enabled digital platforms, the relationship between law and its delivery — previously mediated by the discretion of street-level bureaucrats — will change. As Buffat argues, the impact of this might be ambiguous, and might not necessarily mean a diminution in the affordance of localised discretion.¹⁴³ For example, citizens might in some cases be empowered by access to (digitised) information and resources that were previously unavailable to them, potentially helping to deliver the promise of greater access to legal materials and to justice.¹⁴⁴ At the same time, administrators might opt to exercise discretion *more* readily, in response to the imposition of automated determinations that fail to capture the complexity of the situations they are applied in and to.¹⁴⁵

Empirical research will be required to determine whether that is true in the contexts where RaC systems are deployed. It does, however, raise the question of whether some RaC systems can be effective, even on their own terms. If they are consistently circumvented in order to achieve desirable or just outcomes, this implies something structural about their value in those contexts. It may create complexities and costs that are unforeseen. It also raises the possibility that — paradoxically — the administration of public services becomes *less* transparent, because administrators are forced to circumvent digital systems of governance more readily and comprehensively than they might otherwise have needed to under text-driven law, shifting practices ‘beneath the radar’ in ways that might resist reasonable oversight.¹⁴⁶

3.5.3.4 Formalisation and the shaping of policy

The question of what is deemed formalizable raises a further reflexive issue, namely that the subset of rules which are deemed susceptible to formalisation might end up being treated differently in practice to those that are not. From a certain perspective this is simply inevitable — the purpose of producing RaC-translated rules is that they can be used subsequently in digital systems, which means those rules will be treated and experienced differently by citizens compared to those that are not so translated.

present problems; they would also cause other problems. In the case of the DHSS [Department of Health and Social Security] it would mean that the client’s needs were being routinised more in accord with the needs of the bureaucracy than the needs of the client.’ (Leith, ‘The Application of AI to Law’ (n 16) 33).

¹⁴² Mears and Howes (n 34) 8.

¹⁴³ Aurélien Buffat, ‘Street-Level Bureaucracy and E-Government’ (2015) 17 *Public Management Review* 149, 153.

¹⁴⁴ *ibid* 156.

¹⁴⁵ *ibid* 156–157.

¹⁴⁶ In ‘The Problem with Law in Books and Law in Computers’ (n 112), Leith argues that orality and non-textual practices already make up a large part of legal practice. To the extent that this is a normal part of the balance of legality, putting trust in RaC to remove discretion might end up tipping the balance too far.

This is *per se* neither good nor bad, but the effects of the divergence between these two paths must be anticipated.

One consideration is the extent to which formalisation might come to be seen as attractive *per se*. There may be pressure to extend the reach of RaC translation into domains where previously the rules were thought not to be amenable to formalisation.¹⁴⁷ This is by no means inevitable, but nor is it inconceivable, as governments look for ever greater reductions in the cost of delivering services.¹⁴⁸

Another consideration is the other side of this same coin, and is perhaps even more problematic, if somewhat extreme. Instead of increasing the envelope of rules that are deemed formalizable to include those that were previously deemed unsuitable for it, the impact goes further upstream in the policy process to frame policymaking from the outset by reference to what is formalizable. This idea is reflected in one of the key findings of the Better Rules report, where it states that '[i]t is difficult to produce machine consumable rules if the policy and legislation has not been developed with this output in mind.'¹⁴⁹ There is an implication here that at least some policy and legislation *should* be developed with RaC in mind. If there is a push (however subtle or unintended) toward producing policy and legislative rules that are amenable to formalisation, a significant risk is that policy initiatives that are *not* so susceptible will be deprioritised or even ignored because they do not fit within the adopted RaC approach. At this point, the concern is that areas deemed unsusceptible to formalisation fall off the radar of policymakers because they cannot find a way to make them 'work' within the RaC paradigm. The implicit pursuit of RaC compatibility may inadvertently constrain the legislator to express only those ideas that the formalisation can handle.¹⁵⁰

Shifting further towards formalisation might necessitate the implicit adoption of utilitarian frame that views legal rules as technocratic tools of compliance. While that compliance might be easier to achieve, the cost might be too much constraint on what policies the legislator can express in the code-driven medium. A vision of greater compliance might appear beguiling, but the Rule of Law and legality are about more than that: respect for autonomy and individual dignity necessarily come at the cost of some measure of certainty about whether and how we will comply.¹⁵¹

Thus, even where the RaC translations do not have legal effect themselves, adopting the approach might nevertheless have an effect on text-driven law 'by the back door', by changing the ethos of the processes by which the latter is produced. This would constitute an effect on legal effect of the first kind, to which we can now turn.

¹⁴⁷ Cf. Luca Arnaboldi and others, 'Formalising Criminal Law in Catala', *Programming Languages and the Law 2023 (ProLaLa)* (ACM SIGPLAN) <<https://popl23.sigplan.org/details/prolala-2023-papers/2/Formalising-Criminal-Law-in-Catala>> accessed 12 November 2023.

¹⁴⁸ It is also the explicit vision of some proponents of computational law. See for example Michael Genesereth, 'Computational Law: The Cop in the Backseat' [2015] CodeX - The Stanford Center for Legal Informatics 1.

¹⁴⁹ 'Better Rules for Government - Discovery Report' (n 71) 4.

¹⁵⁰ As Meessen puts it, the decision to formalise "takes away the ability of the legislator to be as expressive about their intentions with the law and burdens them with the responsibility of correctly expressing ideas in a formal language." See PN Meessen, 'On Normative Arrows and Comparing Tax Automation Systems', *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (Association for Computing Machinery 2023) 3 <<https://dl.acm.org/doi/10.1145/3594536.3595160>> accessed 28 October 2023.

¹⁵¹ Waldron (n 16) 19. See also Bańkowski and Schafer (n 16).

3.5.4 RaC and the effect on legal effect

In this final section we can turn to a central element of the broader COHUBICOL analysis: *the effect on legal effect*. By this we refer to two forms of potential impact. First, the capacity or risk that a technology will alter the process of attributing legal effect — that is, will alter the process by which performative speech acts create new institutional facts in specific cases, for example by shaping the legal resources that appear in search results and therefore inform the preparation of a case, or by affecting the text of an AI-generated brief according to the statistical distribution contained in the large language model.¹⁵² There, the nature of the court judgment's legal effect is not changed; it still comes about by declarative speech act, recorded in natural language and susceptible to further interpretation and appeal/overturning/distinguishment by future courts. The means by which the legal effect is arrived at do change (e.g. via the introduction of statistical legal search), but the result is not different in kind. There are concerns here, of course, which were discussed above in Chapter 2, but they are of a different kind to the second form of legal effect.

The second notion of effect on legal effect is where the underlying concept is itself altered, shifting it away from institutionality as its fundamental characteristic, with all that is implied by that. This is more fundamental to the mode of existence of law, because as we saw in the discussion of normativity above, what makes an effect qualify as *legal* is its compatibility with and co-dependency on anchoring practices of attribution, interpretation, contestation and adjudication, and the legal-institutional artefacts that flow from them (norms, rights, duties, personhood, etc.). Taken together, these are ultimately what afford the law's mode of existence, and its capacity to provide protection.

3.5.4.1 The mode of existence of code-driven law

The mode of existence of text-driven law is of institutional facts brought into being by performative speech acts that accord with the requirements laid down in positive law.¹⁵³ Each of these elements affords a specific aspect of legal protection: the institutional fact is not physical; it cannot enforce itself in the way that e.g. a speed bump in the road can. The validity of the speech act is contingent on it properly reflecting the requirements laid down, textually, in the relevant positive law (legislative and judicial). Those requirements will cover a multitude of factors: who is capable of performing the act, under what circumstances, in which particular jurisdiction, and when. The essential contingency of textual meaning implies that the door is always open to contestation of the institutional fact, on the basis of differing interpretations of what those requirements are or what they ought to have meant in a particular circumstance.¹⁵⁴

If natural language is replaced by code as the basic 'dependency' underpinning the mode of existence of the law, then so too will these fundamental building blocks change.¹⁵⁵ This may or may not be a good

¹⁵² For a discussion of the normative implications of this kind of 'third voice' being introduced into legal practice, see Diver and McBride (n 57).

¹⁵³ The standard example is a marriage, which cannot be pointed at but is nevertheless very real. Its creation is a speech act performed by an authorised celebrant, which to be successful must be in accordance with the relevant legal provisions, while also always being subject to potential contestation on the basis of those provisions and other rules and principle that have a bearing on their meaning.

¹⁵⁴ Mireille Hildebrandt, 'The Adaptive Nature of Text-Driven Law' (2021) 1 *Journal of Cross-disciplinary Research in Computational Law*.

¹⁵⁵ Diver and others (n 3) ch 4.

thing, but it will undeniably represent a profound shift, the implications of which are difficult fully to anticipate.

At the level of fundamental legal concepts, the representation in code of the individual institutional facts that are instances of those concepts becomes necessary in order to facilitate automated execution. There must be a translation between the legal world and the code world. Code is not found, but must be created, and so there must be a conscious translation between the two domains. When this happens, the programmer's understanding will necessarily mediate the translation and frame the representation of legal normativity via computationally tractable methods. How they make sense of legal concepts of right, duty, personhood, etc will impact on the approach they take to representing those concepts within the programming tools that are available to them.¹⁵⁶

The further to the right of the normative spectrum a code-driven system sits, the greater the potential gap that must be bridged between a legal concept and its computational representation. The more we move away from the text-driven mode of existence of law, the smaller the role played by attribution and interpretation and the flexibility they afford in our shared understanding of what the constellation of legal effect is at any given moment. Rights become permissions and access control lists, duties become predefined paths that channel behaviour within the interface of the system, contracts and legislation become — as we have seen — 'if this, then that' algorithms, and legal personhood becomes an instance of a user within a predefined role granted a certain range of permissions. Complex interplays between these elements will arise, but these too will be defined in code, rather than natural language. This is inevitable, because those are central elements of modern software of any complexity (albeit not of most contemporary RaC systems towards the left of the normative spectrum). What happens to legal concepts and the relationships between them when they are cast in code? Legal rights and duties are not directly analogous to computational permissions, even if superficially they might sometimes appear that way. What a human can do as a user within a software system is not the same as what a legal person can do within the legal system, however detailed the modelling.¹⁵⁷ What goes on in the latter is of a different category to what goes on in the former. Once we step into the computational environment, and start to 'do law' there (either explicitly or de facto), we are forced to (re)frame legal concepts using the tools and representations that are available in that environment.

3.5.4.1.1 The ecology of code versus the ecology of law

Looked at from an affordance perspective, we can think in terms of law as an 'ecology of practice' whose elements constitute its nature:¹⁵⁸ personhood, rights, duties, and norms have the character they have because of the 'habitat' within which they exist: a shared social world of institutional facts.¹⁵⁹ To

¹⁵⁶ Hohfeld's seminal theoretical analysis of legal relationships has, for example, been a frequent subject of logical formalisation attempts. See e.g. Réka Markovich, 'Understanding Hohfeld and Formalizing Legal Rights: The Hohfeldian Conceptions and Their Conditional Consequences' (2020) 108 *Studia Logica* 129. For Hohfeld's original analysis see Wesley Newcomb Hohfeld, 'Fundamental Legal Conceptions as Applied in Judicial Reasoning' (1913) 23 *Yale Law Journal* 16.

¹⁵⁷ 'It is not simply that business or law are more complex than computer configuration (that, say, they have more 'variables'). Rather it is that they are qualitatively different.' (Leith, 'The Application of AI to Law' (n 16) 34).

¹⁵⁸ Isabelle Stengers, 'Introductory Notes on an Ecology of Practices' (2005) 11 *Cultural Studies Review* 183 (referring to an ecology of practice as 'a tool for thinking through what is happening').

¹⁵⁹ Bruno Latour, *An Inquiry into Modes of Existence: An Anthropology of the Moderns* (Harvard University Press 2013) ch 13; Neil MacCormick, *Institutions of Law: An Essay in Legal Theory* (Oxford University Press 2007).

paraphrase Stengers, law has ‘no identity of a practice independent of its environment... the very way we define, or address, [legal] practice is part of the surroundings which produces its ethos’.¹⁶⁰

Addressing the practice of law in code-driven terms risks changing its ethos and its identity. If the scope of the role played by the legal-institutional environment is reduced in favour of computation, the specifically legal nature of the concepts it can afford, including personhood, rights and duties, will also be curtailed. The legal nature of those concepts cannot be represented in code, no matter how complex the formalism: computational representations by definition exist within a computational environment (even if it is possible for them to be interpreted from outside, this is an external perspective).¹⁶¹ This is necessarily separate from the environment upon which law-as-we-know-it relies, namely (i) the materiality of text and the baseline set of affordances it provides that make law-as-we-know-it possible,¹⁶² and (ii) the legal-institutional ecology that is built on it, consisting of institutional facts and the affordances internal to the law that facilitate legal operations between legal subjects on a constant basis.

While computation, like law, relies on material facts (computers, keyboards, code, developments paradigms etcetera are all ‘in the world’),¹⁶³ the ecology that these establish is fundamentally different from the ecology that law relies on. The text-driven legal ecology affords legal institutionality, whereas the ecology of computation has a categorically different set of affordances built around a specific notion of information processing. Computation results in a set of data structures and operations that can be performed on those structures that are defined within a very specific and limited paradigm of informational representation, albeit one that is very powerful.¹⁶⁴ If we seek to embody legal concepts *directly and substantially* within that ecology, rather than relying on it only as a useful tool for representing the *form* of those concepts (e.g. via digitised documents), we change their mode of existence: we see only the *umwelt* of bits, rather than the *welt* of institutional facts.¹⁶⁵ Digital infrastructure can provide the tools to facilitate the latter, but it cannot replace them directly. Doing so might be an explicit choice in more radical forms of RaC, or it might be an unintended consequence of seeking to fully represent legal normativity via technological means.

If natural language and the legal-institutional environment it affords are sidelined, the current concept of legal effect is necessarily also sidelined. With it go its affordances of attribution (declaring a legal state of affairs and thus bringing it into being), flexibility of interpretation (to determine what it means in a particular circumstance for a given person), contestability (to allow a formal mechanism to challenge

¹⁶⁰ Stengers (n 158) 187.

¹⁶¹ AJ Wells, ‘Gibson’s Affordances and Turing’s Theory of Computation’ (2002) 14 *Ecological Psychology* 140, 171.

¹⁶² The distinction and interplay between what affords law, and what law affords, is important here. See Diver, ‘Law as a User’ (n 34) 22ff.

¹⁶³ Even Turing’s abstract machine was ‘ecological’, in the sense that it was built around a very limited set of physical operations performed on a physical tape (Wells (n 161) 171). This idea is vividly portrayed in chapter 13 of Liu Cixin’s science fiction novel *The Three-Body Problem* (Ken Liu tr, Head of Zeus 2016).

¹⁶⁴ Shannon (n 104). See also Joseph Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation* (Freeman 1976) ch 3; Hildebrandt, ‘Law as Information in the Era of Data-Driven Agency’ (n 104).

¹⁶⁵ See Laurence Diver, ‘3.5.5.3 Protecting the legal subject by protecting the mode of existence’ in Diver and others (n 3) 91.

that meaning), and adjudicative closure (to enforce the decision in a way that the community as a whole can understand and respect, even if they disagree).

3.6 Conclusion: treading the line

The shift to computation need not be as wholesale as was just described. Indeed, many RaC systems are not intended to supplant legal effect; they can be deployed as tools to be used in service of text-driven normativity, impacting legal effect in the first sense mentioned above, that is by shaping how legal effects are attributed in specific circumstances, rather than changing the underlying nature of the concept. That influence can be positive, in terms of legal protection and the Rule of Law, as for example where access to relevant legal information and interpretative materials strengthens understanding of the landscape of legal norms, and empowers people to avail themselves of the law's protective capacity by enabling them to make stronger and more creative arguments in defence of their rights. Delineating the proper role of computation is essential, to ensure it complements and strengthens *legal* normativity, rather than sidelining it or converting it on-the-fly into a technological normativity that lacks the protective affordances of text-driven law.