Comments

on

Cracking The Code:

Rulemaking For Humans And Machines
(August 2020 draft)

Comments on the draft OECD White Paper on Rules as Code,
submitted on 27 August 2020 to the authors

October 31, 2020

LawTech La Trobe Research Group

Report N° LawTech/02/Oct-2020

Authors:

Jeffrey Barnes	`j.barnes@latrobe.edu.au`
Pompeu Casanovas	`p.casanovasromeu@latrobe.edu.au`
Louis de Koker	`l.dekoker@latrobe.edu.au`
Guido Governatori	`g.governatori@data61.csiro.au`
Mustafa Hashmi	`mhashmi@ltu.edu.au`
Ho-Pun Lam	`brian.lam@csiro.au`
John Zelenikow	`j.zeleznikow@latrobe.edu.au`

Compiled By:

Mustafa Hashmi	`m.hashmi@latrobe.edu.au`
Pompeu Casanovas	`p.casanovasromeu@latrobe.edu.au`

CONTENTS

PREFACE

Cracking the Code: Rulemaking for Humans and Machines is an OECD Working Paper authored by James Mohun and Alex Roberts and released on 14 October 2020. This Working Paper on Public Governance, produced by the OECD Observatory of Public Sector Innovation, is available in English and French. A draft version was publicly available on the web for comments. In August 2020 the La Trobe LawTech team reached out to the authors and provided some comments on the draft Working Paper. Rules as Code is gaining momentum, increasing the need for constructive critical engagement. To support the broader discussions around Rules as Code, we are publishing our August 2020 Comments. These are provided in their original form, with small technical and language edits for wider publication.

1 An AI or 'Rules as Code' Project: The Implications of Statutory Interpretation

Jeffrey Barnes

Summary

An AI or 'rules as code' (Mohun and Roberts, 2020) project is not credible without an adequate examination, and taking into account, of the implications of statutory interpretation in the legal system. The questions that should be asked include:

- Why is interpretation of the law a consideration in an AI or 'rules as code' project?

- How important is interpretation of the law in such a project?

- What implications does the proper consideration of interpretation have on a rules as code project, in particular on the status of the code and on its possible role more generally in the legal system?

1.1 Preliminary

Although my references tend to draw on Australian and United Kingdom law, the principles of interpretation of legislation have much in common throughout the common law and civil law worlds (MacCormick and Summers, 1991).

1.2 Why is interpretation of the law a consideration in an AI or 'rules as code' project?

Law as made is 'incomplete' (Pistor and Xu, 2003). This is a well-established thesis of many courts, parliamentary counsel, and scholars. For instance, Bennion, a former United Kingdom parliamentary counsel and, at the time of writing a leading scholar on legislation, wrote:

> A law text, even if it is an entire Act, is far from being the whole story. Every Act is incomplete in itself. Law is a palimpsest or multiple imprint surface. An individual law text needs to be considered in context. No one law text stands alone. It always needs to be read alongside many other law texts, and this cannot be achieved by unaided non-lawyers (Bennion, 2007).

Interpretation is part of the process of making law. What Donaldson J said of judges is true of all interpreters—in the executive as well as in the legal profession:

> The duty of the courts is to ascertain and give effect to the will of Parliament as expressed in its enactments. In the performance of this duty the judges do not act as computers into which are fed the statutes and the rules for the construction of statutes and from whom issue forth the mathematically correct answer. The interpretation of statutes is a craft as much as

a science and the judges, as craftsmen, select and apply the appropriate rules as the tools of their trade. They are not legislators, but finishers, refiners and polishers of legislation which comes to them in a state requiring varying degrees of further processing.
(*Corocraft Ltd v Pan America Airways Inc [1969] 1 QB 616, 638*).

A statute does not exist on its own as law. Law is a 'network' (Barnes, 2013, p.49). To work out the law requires much more than a linguistic analysis of the structure of a legislative provision. Here are authorities from Australia and the United Kingdom:

> The meaning of a statutory text is also informed, and reinforced, by the need for the courts to apply the text each time, not in isolation, but as part of the totality of the common law and statute law as it then exists (Gageler, 2011, 1-2)

> . . . any statute must of course be looked at in the light of the general law of the country. Parliament in its wisdom in passing an Act must be taken to know the general law.
> (*Fisher v Bell [1961] 1 QB 394, 399*)

The need for interpretation in the sense of resolving ambiguities (disputes over the scope of a statutory provision) is ever present by reason of 'sources of doubt'. Interpretation is not work created by judges and the legal profession; rather it arises because of sources of doubt. They 'are the events, decisions and other factors from which doubt may arise about the meaning and application of a legislative provision in a particular case' (Barnes, 2008, p.120). There is substantial learning on sources of doubt (Twining and Miers (2010, ch. 6);(Bennion, 1990, chs 15-19); Barnes (2008); Schane (2006, ch. 1)) which law-makers and policy makers ignore at their peril. This is because of the qualities of sources of doubt. They are: numerous and diverse; not restricted to events and processes taking place before the Act was passed; ineradicable; impacted by the circumstances of each case; and able to be minimised but not eliminated by the use of plan language drafting techniques (Barnes, 2010).

1.3 How important is interpretation of the law for an AI or rules as code project?

The scope for interpretation should not be overstated. But neither should it be underestimated. Some misconceptions should be pointed out. Estimates vary as to the extent to which legal meaning of legislation varies from its ordinary meaning (the meaning that is likely to be programmed in an AI project[1]). The High Court of Australia has said that 'ordinarily' legal meaning corresponds with the grammatical meaning: *Project Blue Sky Inc v Australian Broadcasting Authority (1998) 194 CLR 355, 384 [78]*. See also authorities in (Barnes, 2008, p.121-6) . However, in a paper delivered to the Australian Academy of Humanities, Professor Colin Howard thought that the peculiar characteristics of legal language meant that there was only a 'superficially close connection with the ordinary language of everyday' (Howard, 1993, p.29). It is true that courts only have a limited role in the legal system. But it is a misconception to equate statutory interpretation with the work of the courts. The work of the courts is a relatively small proportion of the work of interpreters, most of whom are located in the executive branch of government and in the legal profession.

---

[1]Programmers are likely to suffer the same inherent biases of drafters. As a general rule, a drafter does not rely on maxims of interpretation (Quiggin, 2014, p.88); cf (Moran, 1991, p.109). The drafter relies on the 'natural' reading (Quiggin, 2014, p.88) or that words mean what they say (Penfold, 2006, p.488); Quiggin (2014, p.88).

1.4 What implications does the proper consideration of interpretation have on a Rules as Code Project, in particular on the status of the Code and on its possible role more generally in the legal system?

This question is difficult to answer in the abstract. It is not because of the uncertainties of interpretation; it is because of the uncertainties of what is attempted in an AI or 'rules as code' project. This uncertainty has two dimensions:

- What status is proposed or attempted for the rules as code?

- What area of law is rules as code applied to?

Part of the likely success of rules as code depends on the status it is accorded. As regards the first dimension, let me consider 3 possibilities:

First, rules as code purports to have authoritative status. It is impossible to see how such a project is feasible. Questions which arise include

- How would it be constitutional? If the governing law purported to deprive the courts of adjudicating for themselves, it would breach the separation of powers[2].

- Assuming it were constitutional, how would the governing law be interpreted in the light of: apparent conflict with other provisions in the same Act, unanticipated developments, programming errors, and a host of other potential sources of doubt?

Second, could it operate as an 'official' document in the manner of, as has been suggested, an Explanatory Memorandum? If the rules as code were attempted after the law was passed, it would be regarded as self-serving and not an aid to interpretation: *Barnes v State of Victoria [2015] VSCA 343, [46]*. However in fairness, some proponents propose an AI project in tandem with the preparation of legislation. The better analogy is not with an Explanatory Memorandum but with an example. An Explanatory Memorandum (if useful) gives explanatory background to a Bill. Algorithms are different; they purport to give answers. An example is a close analogy with algorithms because an example purports to give an answer to stated facts. There is nothing unconstitutional about a non-binding example. Examples are common in plain language drafting (Barnes, 2004). However, at common law examples do not take priority over a rule. If an example is inconsistent with the relevant rule, the rule prevails: *Ariffin v Gark [1916] 2 AC 575, 580, 581.* Also, to be of assistance, the algorithm's answer would need to be reasonably attributable to Parliament. It could not be attributed if the analysis was not open and apparent to lay members of Parliament.

Third, rules as code purports to be an administrative aid to decision-making. There is some support in the literature for rules as codes for 'mundane' administrative work (Derkley, 2020, p.14); what is also called in the literature 'easy cases'. You can get 'efficiency gains' by doing 'things in high volume at high speed' (Derkley, 2020, p.14), quoting Katie Miller). But here, because it lacks status as an aid, it would be misleading to describe such technical support as 'official'. It is more an administrative support.

---

[2]In Australia and other countries where separation of powers is a higher law doctrine, this would spell invalidity: *Taylor v The Owners — Strata Plan No 11564 (2014) 253 CLR 531, 549 [40].*

1.5   What overall conclusions can be drawn?

Casting his net wider than I have done, Bateman says that 'many proposals to automate statutory powers are likely to be legally faulty' (Bateman, 2020, p.520). I would agree with this. A rules as code project has strong echoes of, and can therefore learn from, the plain language movement. Like the plain language movement, the rules as code proponents are propelled by well-intentioned goals of making the law accessible. But the plain language movement has failed to achieve the strong goals of the movement and there is no prospect of it doing so. It only works 'in a limited sense' (Assy, 2011; Barnes, 2013; Bennion, 2007, p.270). While plain language drafting techniques have been of value for the democratic process and have brought about improvements to the legal system (Barnes, 2013, p.273), they have been unable to achieve legal certainty, that is, quell sources of doubt that have to be resolved with legal expertise' (Barnes, 2013, p.273).

2 Rules as Code: The need for an impact assessment to inform application

Louis de Koker

Chapter 8 of the draft OECD report addresses the type of rules that lend themselves to coding (Mohun and Roberts, 2020). It is submitted this question goes beyond the type of rules and that, in practice, a more holistic approach will be beneficial when the application of Rules as Code is considered. While the general vision appears to be to code all types of legal rules, it is submitted that, in practice, most of the applications will be in the context of what Barnes describe as support for administrative work. This means in essence that it will operate at a level of business rules, unpacking legal rules (Barnes, 2020). Examples of these are the welfare and tax calculators and web-based guidance to government services that are often encountered. This approach should consider the benefits of coding certain rules, the risks of coding them and the management of foreseeable consequences. This can best be done by undertaking a transparent and public impact assessment. In this brief submission I will touch on only two aspects that should be assessed: Efficiency and liability. In addition, however aspects such as capacity, methodology, ethics and costs should also be considered.

2.1 Efficiency

Rules as Code promises efficiency: Instead of numerous companies coding their own interpretation of the rule, a coded version issued by the rule maker is published that all users can employ (Barnes, 2020). This promise is more likely to be true where the rules are simple and lend themselves to coding that will not be subject to any interpretational challenges and will hold true in all cases, foreseen and unforeseen. Human interpretation may be inefficient but interpretation, as Barnes argues, is indispensable as it allows for a measure of flexibility to determine whether and how the rules should be applied to the facts, especially in novel circumstances. Coding introduces a measure of rigidity in the application of the rule. When exceptions to the rule are coded, the exceptions are limited to those that are identified and anticipated by the coders. It is likely that novel circumstances will be encountered that are not envisaged in the coded version and that remedial action will be required to provide a fair and just solution. That remedial action introduces a measure of inefficiency that should be discounted when the overall efficiency of coding of the rule is assessed.

2.2 Liability for errors and remediation

The coded version of the rule may be challenged and found to be inconsistent with the natural language version of the rule. It is also foreseeable that the interpretation of the natural language text of the rule may be challenged in a court of law and may result in a different interpretation to

the coded interpretation. The report itself supports "appealability" of the coded version, allowing for the code to change, potentially after having been operationalized (Mohun and Roberts, 2020).

Coded rules lend themselves to large-scale application, whether to provide information and guidance (for example regarding entitlements and access to government services) or to support automated processes, including automated decision-taking. The consequence of such an application of an incorrect interpretation should be considered upfront, as well as questions of liability for any negligent mistake and for the costs of remediation.

Australia's so-called "Robodebt" scandal is not an example of Rules of Code but does illustrate the impact of the large-scale automated application of an incorrect principle to determine and reclaim overpayment of social assistance (Commonwealth of Australia, 2017; Keyzer, 2017). Debts were assessed and asserted on the basis of overpayments suggested by data-matched estimates of averaged fortnightly earnings. In December 2016 the government began to send letters of demand requiring immediate repayment and purported to shift the onus of proof that no debt existed, to recipients. After a legal challenge, the government conceded in December 2019 that this approach to debt assessment was not valid[1]. It is reported that the Australian government will be forced to refund more than 400,000 welfare debts worth about A\$ 550 million that were wrongly issued to hundreds of thousands of Australians (Henriques-Gomes, 2020).

Ideally the agency concerned should formally adopt responsibility for any errors in the coding or any corrective action that is required. The Australian Taxation Office (ATO) provides a good example if an agency that did not accept the consequences of the technology it adopted.

The ATO adopted template letter technology to enable its officials to issue letters to tax payers. On Monday 8 December 2014 the ATO sent a letter to the taxpayer bearing the signature block of the first Deputy Commissioner, headed "Payment arrangement for your Income Tax Account debt." The letter thanked the taxpayer (Mr Pintarich) for a recent promise to pay outstanding amounts and stated[2]:

> We agree to accept a lump sum payment of A\$ 839,115.43 on or by 30 January 2015. This payout figure is inclusive of an estimated general interest charge (GIC) amount calculated to 30 January 2015.

The taxpayer made payment in full of the lump sum referred to in that letter on 30 January 2015. The ATO, however, proceeded to claim the full GIC (and amount of approximately A\$ 335,000), arguing that the statement in the letter was not intended and that no decision was taken regarding the remission of the GIC (own emphasis):11

> Mr Celantano (ATO official) said that he had caused the letter to issue but was unable to explain how the sentence in the second extracted paragraph (regarding GIC) had come to be included. He had "keyed in" certain information into a computer-based "template bulk issue letter". This process had generated the document. He had not read the letter before it was despatched. He deposed that what was said in the first two of the paragraphs, extracted above, did "not accord with the conversations" he had had with Mr Pintarich and Mr Smith. Mr Celantano said that he had not, at any time, made any decision, under s 8AAG, to remit any GIC, owing by Mr Pintarich.

While the majority decision in the Federal Court sympathised by the tax payer, the Court agreed with the ATO that no decision was taken in relation to the GIC:12

> If the natural reading of the December 2014 letter is as set out above, it would follow that the letter communicated that a decision had been made to remit all GIC payable by the

---

[1] *Deanna Amato v The Commonwealth of Australia VID/2019/611-0* `https://www.comcourts.gov.au/file/Federal/P/VID611/2019/3859485/event/30114114/document/1513665`

[2] *Pintarich v Deputy Commissioner of Taxation [2018] FCAFC 79, par 3.*

taxpayer save for the relatively small amount of GIC covered by the lump sum payment amount referred to in the letter, if the taxpayer paid the lump sum on or before the specified date. However, we do not consider that this resolves the question whether the Deputy Commissioner made such a decision. In order for there to be a decision to remit GIC under s 8AAG of the TA Act, we consider that there needs to be both a mental process of reaching a conclusion and an objective manifestation of that conclusion. In the present case, on the basis of the findings of the primary judge (which are not challenged on appeal) there was no mental process of reaching a conclusion.[3]

The Pintarich case is not an example of Rules as Code. It is however an example of an agency that adopted technology but refused to be bound by a mistake made when the technology was used. What would the attitude be if the agency discovers that the coding was incorrect and led to errors detrimental to the agency and the government? What would be the processes if the error was to the detriment of thousands of citizens? Upfront and transparent answers to questions such as these should be part of the processes to consider which rules are to be coded and by whom.

---

[3] *Pintarich v Deputy Commissioner of Taxation [2018] FCAFC 79 par 151*: "We accept that there may be some perceived unfairness in the circumstances of this case. As indicated above, in our opinion, the natural reading of the December 2014 letter, in the context in which it was written, is that the Deputy Commissioner agreed to accept a payment of the lump sum amount on or before 30 January 2015 in full discharge of the taxpayer primary tax and GIC liabilities as set out in the statement of account dated 10 November 2014. The taxpayer duly paid the lump sum amount by the due date. It may be perceived as unfair if the Deputy Commissioner is able to go back on what was conveyed by the letter. However, the issue to be determined is whether or not, on or about 8 December 2014, a decision was made to remit GIC. As explained above, we consider that, for such a decision to be made, there needs to be both a conclusion reached on the application to remit as well as an overt act. In the present case, neither Mr Celantano nor anybody else had reached a conclusion on the application for remission at the relevant time."

3 Comments on Cracking the Code. A short note on the OECD Working Paper Draft on Rules as Code

Pompeu Casanovas

3.1 Introduction

This is a comment on the Draft paper launched for public consultation by the OECD's Observatory of Public Sector Innovation (OPSI) within the Open and Innovative Government Division of the Public Governance Directorate in June 2020. The OECD (2017) has also published a useful report on algorithms, business and policy:

> Although few would dispute the great benefits offered by algorithms, especially in terms of improved automation, efficiency and quality. [. . .] there are questions about the extent to which human decision-making will be supported (or even replaced in certain cases) by machines and the implications of the automation of decision-making processes for competition. [. . .] researchers must not create something which cannot be controlled [emphasis added].

The objectives of Rule as Code are presented by the authors of the Draft as follows:

> Rule as Code (RaC) is an exciting new concept that rethinks one of the core functions of governments: rulemaking. It proposes that governments create an authoritative version of rules in software code which allow rules to be understood and actioned by computer systems in a consistent way [emphasis added]. This challenges the long-established processes of government rulemaking and could transform both policy and public service delivery. It envisions and helps support a truly digital government, with rules created as a digital product and service, rather than having them incorporated into digital processes after the fact. It creates the conditions for a government that can be more agile, more responsive and more innovative in navigating and shaping an unpredictable operating environment.

To briefly comment their findings from a Law and Technology perspective, I will keep separated three different dimensions of law. All three are necessary to define law as knowledge, and more precisely, to apply algorithms and formal rules to law: (a) law as data; (b) law as system; and (c) law as theory. All three dimensions also assume (i) that legal contents—expressed in natural language—can be comprehensively represented into some formal language; (ii) that legal contents can be used for many purposes. For instance, it is not the same to structure law as data to be stored, archived and downloaded from a library than building an automated code (a) to apply legal rules to a population, (b) define and manage citizens' rights and duties, (c) and incentivise, punish or fine them through negative and positive sanctions. In the first case, law is understood as a linguistic resource and treated as information. In the second one, law is defined under the rule of law and refers to a system of norms or rules created and enforced through statutory and/or case-based law by judges, government agencies and Legal Enforcement Agents (LEA).

At the present time these three dimensions are related on the Internet of Things (IoT) through the architecture, languages, recommendations and standards of the Semantic Web (SW)[1]: "web of data (or data web) that can be processed by machines—that is, one in which much of the meaning is machine-readable" (Wikipedia) via the World Wide Web (WWW). Information system where documents and other web resources are identified by Uniform Resource Identifiers (URIs, such as `https://example.com/`), which may be interlinked by hypertext, and are accessible over the Internet. The WWW Consortium (W3C)[2] is "an international community where Member organizations, a full-time staff, and the public work together to develop Web standards". This was created following Tim Berners-Lee's particular I have a dream (Berners-Lee and Fischetti (1999); Berners-Lee et al. (2001)). In the last few years, the SW has evolved into a Web of Linked Data[3] As also described by W3C:

> The Semantic Web is a Web of Data — of dates and titles and part numbers and chemical properties and any other data one might conceive of. The collection of Semantic Web technologies (RDF[4], OWL[5], SKOS[6], SPARQL[7], etc.) provides an environment where application can query that data, draw inferences using vocabularies, etc.

3.2 Law as Data

Law as data takes all regulatory components—rules, norms, schemes, directives, principles, values, plans and strategies, to mention just a few—as information for the purposes of legal data management, definition, storage, transfer, modification, retrieval and usage. Already in the 1990s, this broad notion has been linked to databases, lexicon vocabularies, data models and language resources. Language resources are defined as *"pieces of data containing linguistic information in machine readable form"* (Marin-Chozas et al., 2019, p. 170): (i) glossaries and terminologies, (ii) lexical databases; (iii) dictionaries, (iv) thesauri (hierarchical controlled vocabularies). Some terminologies in the legal field are quite specialised, e.g. Ontolex-Lemon (OL) is a database of concepts and terminological data related to the copyright and published in [TBX format – ISO 30042] and also in RDF, suitable for establishing links with other resources such as DBpedia (Rodriguez-Doncel et al., 2018). OL has extracted some terms from IATE[8], the big interactive terminology database of the EU (Moreno-Schneider et al., 2020a).

The first layer of the (Legal Semantic Web (URIs, XML, RDF) have been developed for statutory and case-based law. The European Legislation Identifier (ELI[9]) and the European Case Law Identifier (ECLI[10]) format, so that it can be accessed, exchanged and reused across borders". They are based on the assignments of URIs, description of metadata and sharing of metadata in machine-readable format.

There is a roadmap for law as data (Rodriguez-Doncel, 2019.). It is pointing out the building of legal knowledge graphs (Moreno-Schneider et al., 2020b): (i) to ordering and managing literally billions of legal data and (ii) and to allowing the linkage of thousands of sources (Moreno-Schneider et al., 2020a). This is a field that is progressing at an extraordinary path. Due to Machine Learning (ML) and AI, NLP techniques are entering into a new stage. Recent work by Brown

---

[1] Wikipedia. Semantic Web. `https://en.wikipedia.org/wiki/Semantic_Web`

[2] World Wide Web Consortium. About W3C. `https://www.w3.org/Consortium/`

[3] WWW Consortium. Linked Data. `https://www.w3.org/standards/semanticweb/data`

[4] Resource Description Framework `https://www.w3.org/RDF/`

[5] Web Ontology Language `https://www.w3.org/OWL/`

[6] Simple Knowledge Organization System `https://www.w3.org/2004/02/skos/`

[7] `https://www.w3.org/2009/sparql/wiki/Main_Page`

[8] IATE. Interactive Terminology for Europe. `https://iate.europa.eu/home`

[9] `https://eur-lex.europa.eu/eli-register/about.html`

[10] `https://e-justice.europa.eu/content_european_case_law_identifier_ecli`

et al. (2020) has demonstrated that scaling up language models improves their performance. GPT-3[11], an autoregressive language model with 175 billion parameters, can generate samples of new articles—including lawsuits and legal documents—which human evaluators can barely distinguish from documents written by humans. This raises ethical concerns about its possible social, political and economic impact.

3.3 Law as a system

As said, in the past two decades SW languages have been developed for the legal domain. Meta-lex, Open XML Interchange Format for Legal and Legislative Resources, has been implemented in the UK and in the Netherlands; CHLexML[12] in Switzerland, NormeInRete[13] in Italy, Akoma Ntoso[14] [AKN] in USA, etc. From 2010 onwards, Legal Web Services are offering tools with semantic parsing and data analytics. There is a convergence between the emergence of the so-called LawTech, FinTech and SupTech companies, and the development of (i) legal analytics (machine learning and deep learning, data mining and advanced statistics); (ii) legal argumentation (AI & Law); (iii) the increasing use of W3C language recommendations and standards—linked data (RDF), ontologies (RDF, OWL), query languages (SKOS and SPARQL), inference (OWL, RuleML) and vertical applications (in specific domains such as law, health, e-government etc.), (iv) the parallel development of semantic rule-based languages in other standardisation organisations (e.g. LegalXML and legalRuleML in OASIS).

SW languages can be used to generate the automated selection and aggregation of relevant information (Van Opijnen and C., 2017) to perform valid legal acts (Casanovas et al., 2016). This is a challenge, because legal data management is no longer just handled as information retrieval and data sharing, but it becomes a mechanism to carry out legal operations. Thus, the overall focus when dealing with legal data is shifting (i) from the need to publish legal documents to an effort to produce complex applications providing legal reasoning, arguments and solutions; (ii) from the organisation of documents (libraries) to the extraction, production and transformation of legal meaning.

It is possible to formally model the law only if: (i) the selection of formal languages is performed after a 'thick' knowledge acquisition process to set the list of technical requirements; (ii) the expressivity of the selected formal languages is previously defined (assessing the range of their scope and limitations), (iii) the field of application is parametrised according to its properties (contracts, tort law, criminal law...), (iv) a variety of interpretations are envisaged in different possible scenarios, (v) internal and external mechanisms are set to monitor and control the automated outcomes within the implementation process.

Even though, lessons learned in twenty years of legal ontology building and, now, in privacy, data protection and security by design or by default, lead to the same result. The overall legal processes of creation, interpretation and implementation of normative lifecycles cannot be fully hardcoded, as natural languages and human behaviour intertwin, mix up, are socially evolving—languages have an "open texture" to be used in multiple situations—and are contextually bounded and determined. Trying to solve specific problems—such as embedding the General Data Protection Regulations (GDPR) provisions into the information processing flow within specific platforms or create reusable ontologies or ontology design patterns (ODP) for the legal domain, brings about semi-formal, tactical, indirect or hybrid solutions (Casanovas et al. (2014); Colesky et al. (2016);

---

[11]Open AI GPT-3 https://openai.com/blog/openai-api/.
[12]https://lawin.org/chlexml/
[13]https://lawin.org/normeinrete/
[14]http://www.akomantoso.org/

Koops and Leenes (2014)) Coding per se was not Berners-Lee's 'dream', as it was not Sergot's and Kowalski's intention when they applied the immanent logic of Prolog to formalise the British National Act (Karagiannis, 2008) On the contrary, they tried to face some unsolved problems and refine the expressivity of formal languages to capture intentions, plans, actions, competences, powers and rights that were written in legal natural language so far.

Thus, the legal process is better described as an institutional process in which coding, legal interpretation, the creation or refinement of algorithms and formal languages, and the pragmatic view of many stakeholders in government, the market, and the political arena (including consumers' and Human Rigths organisations) convene to create specific regulatory models (RM). Models of rules express normative systems with human or artificial agents—the so-called in AI normative Multi-Agent systems, socio-technical systems or socio-technical cognitive systems (Andrighetto et al., 2013). They also try to create the conditions for the coordination of agency and frame its implementation at the same time—i.e. its ecosystem. And there are many methodologies to systematically carry it out with a variety of tools (smart contracts, blockchain...) in different situations and scenarios (taxation, allocation of resources, wellbeing, services of health or age care...). Regorous (Governatori, 2015b), Formal Legal GRL (Rabinia et al., 2020), Legal-Urn and Eunomos (Boella et al., 2014) are examples of such methodological trends to face compliance in the legal domain (Governatori et al., 2016). I will recover this thread later, in my comments on the OECD Draft.

3.4   Law as theory

Computer representation languages and normative and regulatory systems must reflect and encompass a variety of regulatory tools: (i) different types of legal instruments whose use is in itself complex (i.e. non-linear)—EU directives, constitutional principles and state planning, legislation, case-based law, policies, agreements, contracts, standards, ISOs, ethical values..., (ii) the interrelation of fundamental rights and duties coming from the general model of the rule of law according to a variety of legal cultures (common, civil and transnational law); (iii) the provisions of formal and substantive due process of law; (iv) the general principles and values of human rights; (v) ethics.

Prior to perform any modelling, the selection of sources and tools, and the theoretical approach to analyse them must be specified. There also is a plurality of (in)compatible perspectives on the attributes and values of legal norms (validity, efficacity, effectivity, enforceability, etc.). This has been the field of legal theory and analytic jurisprudence for more than 150 years now (if we take into account the classical work by John Austin and Jeremy Bentham at the beginning of 19 c.). There are two important analytical steps: (i) the legal dogmatics of a particular field (i.e. criminal law, torts, contracts...) sets the main concepts—called fundamental legal concepts—that are inducted and fleshed out from legal sources (i.e. property, intellectual property, crime, murder, obligation, etc....), their conditions, schemes and use in legal arguments; (ii) legal theory tries to structure—substantively and procedurally—their relevance, salience, consistency, legitimacy, legality and feasibility under constitutional, political, and philosophical tenets (e.g. the inference of their legal effects). Thus, the development of defeasible, non-monotonic logic, and the so-called non-standard deontic logic, have been important for contemporary AI & Law and Law & Technology developments.

Now, the abstract properties of the languages used, and the logical consistency, semantic coherence, temporality and typology of rules, norms and normative systems have also been modelled through computational means. Researchers have, then, encountered the problems raised by the lack of scalability, interoperability and expressivity of systems, languages and

formal solutions. In one of the most comprehensive papers on interchange legal engineering requirements, Gordon et al. (2009) listed them in an ordered manner— isomorphism, reification (jurisdiction, authority, temporal properties), rule semantics, defeasibility, etc. After examining RuleML[15], SBVR[16], SWRL[17] and RIF they also concluded that there was not (as there is not now) an interchange language which can satisfy at the same time all the requirements they had listed. Law, legal behaviour and natural language, are complex fields.

Does this mean that we should stick with natural language and abandon the attempts to model legal norms? Not at all. No. What we need is more, not less formalisation. It is my contention that these trends are mature enough to be implemented, benchmarked and tested in real settings, under the right conditions in specific areas, i.e. excluding experimentalism, because normative and regulatory systems should be integrated into the democratic legal systems in place, not the other way around. As Presutti et al. (2009) put it once, what we need is favouring the reuse of encoded experiences and good practices, pushing towards an extreme design. And—let me add—pushing towards 'extreme' empirical approaches, tests and cooperation among all stakeholders—end-users, citizens, researchers, and governments as well. This is known as multi-stakeholder (symmetric or asymmetric) governance (Casanovas et al., 2017) for linked open data. It is my contention that, subject to some guarantees, these trends are mature enough to be implemented and tested in collective actions, community developments, government policies, case-based law, and parliamentary legislation.

3.5 Comments the OECD White Paper

For the reasons stated above, the initiative of the OECD is timely and most welcome. I think that its intention is drawing the public attention and fostering a public dialogue about the need of encoding laws and policies (i) to better serve citizens' demands, (ii) to save public money, (iii) and to avoid unnecessary bureaucratic delays and pitfalls in the delivery of government services. I support these objectives. As is also the case with contemporary dictionaries, "machines are a major consumer of government" (Draft, p. 70). Laws and regulations should, and will be, expressed within digital formats in the next future, not only in natural language. The question is how, in what way this will be made possible, and what kind of knowledge will be produced and used to explain and anticipate their impacts. Very likely there will be a lengthy transitional period where we will have the opportunity to learn about their coexistence, frictions and maladjustments to enable us to address them.

Having said that, I have more questions than answers. After reading the paper, and despite the business language and qualifications used by the authors ("new", "exciting", "innovative", "rule consuming" ...), I don't see a clear explanation of who coined the term "rules as code" (RaC), and for what purposes. How did this concept come about? And from where? The authors write:

> In allowing third parties to directly consume an authoritative version of coded government rules, it promises the potential for quicker service delivery, a more consistent application of the rules and greater efficiencies for rule takers. (Draft, p. 7)

> RaC suggests that, if government were to assume the role of rule maker, it could result in stronger alignment between rule intent and implementation. (Draft, p.8)

It is not clear to me what "authoritative" means, but governments under the rule of law cannot set once and for all the meaning of the law. This is a dynamic and evolving process, under democratic

---

[15]http://wiki.ruleml.org/index.php/RuleML_Home
[16]https://www.omg.org/spec/SBVR/About-SBVR/
[17]https://www.w3.org/Submission/SWRL/

controls. Also, citizens and right holders do not "consume" rules. They legally comply (to certain degree) with them and fulfil duties and rights. The authors could consider what 'social contract', 'jurisdiction' and 'separation of powers' mean under constitutional laws. The language they use belongs to computer studies for marketing and business research, but citizens cannot be defined as consumers or employees (these are roles they can assume), and the public space cannot be confused with the market or with corporative scopes and confines.

In the last five years e-government models for the Public Administration have been endorsing corporate rules, organisation models and architectures—i.e. COBIT[18] and TOGAF[19]. In principle, there is nothing wrong with that if the purpose is to increase the efficiency and to reduce risks. However, if this is the case, the authors should clarify first the list of technical, social and legal requirements of the design and explain its rationale, because public principles matter. To put but one example, this is what Mondorf and Wimmer (2016) did to design Pan-European e-government services.

I have found an explanation about the origins of RaC in Waddington (2019):

> "machine-consumable legislation", or "rules as code", which was first raised in New Zealand, spread to New South Wales, and is now being considered in several other Commonwealth countries. Many governments already publish their legislation in a coded form that enables a computer to read identifying features of each provision in the legislation (such as that it is section 19(4)(a)(ii) of the XYZ Act as it was in force on a particular past date). But that coding leaves the computer unable to extract any of the meaning of the provision (beyond searching for words). Many governments also have coded versions of some legislation to do such things as calculate social security entitlements and issue payments, and many commercial firms sell software that performs similar functions for the public. "Machine-consumable legislation" combines these two approaches, so that policy rules would be digitised before and during the legislative drafting process. The resulting coded version could be published (on a site from which computers can access it automatically) in tandem with the enacted legislation.

Thus, the "consumer-rule" view is close to "machine-consumable", referring to information or data being processed by machines. The New Zealand Report includes a definition (New Zealand Government, 2018):

> 'Machine consumable' for the purpose of this report means having particular types of rules available in a code or code-like form that software can understand and interact with, such as a calculation, the eligibility criteria for a benefit (e.g. see the financial assistance eligibility tool for SmartStart, which is powered by a digital rules engine) or automated financial reporting obligations for compliance.

Quoting a Canadian government project[20]:

> Rules as code is the process of translating legislation, regulation and policy from words into code. This involves taking the rules that are written in English/French and converting them into machine readable data and code. This also includes using coding concepts and logic in the initial design of legislative drafting, which should ultimately make legislation clearer and make it more easily machine interpretable.

This is related to compliance (I will come back later to this subject). It does not invalidate my comment in 3.2 and 3.3, but I can see now that RaC is a movement led by developers (M. Jarvis) drafters (M. Waddington), civil servants (Pia Andrews), rather than a research or computer science-driven trend, and aimed at practical objectives. Perhaps this can explain why they do not refer to the bulk of work already done and lessons learned (see sec. 2-4). After thirty

---

[18]https://www.isaca.org/resources/cobit
[19]https://www.opengroup.org/togaf
[20]https://wiki.gccollab.ca/CFR_Experimentation

years of ontologies, it is a bit surprising reading from Waddington (2019)'s slides: "Ontologies v. Interpretation. We have to understand their idea of ontologies [digital folks]." Ontologies are not neutral, and they always entail human interpretation, as it is well-known since the work of Gruber in 1992 on portable ontology specifications. Knowledge is shared, and interpretations are negotiated and discussed in ontology-building processes.

This could perhaps explain as well why this Draft on rulemaking is not based on any analytical or theoretical framework. The paper never mentions the common notions of Semantic Web or Web of Data or Web of Linked Data, Rule of Interchange Format (RIF), World Wide Web (WWW), LegalRuleML, ODR (Open Digital Rights), ODR (Online Dispute Resolution) etc. The OECD Draft is mainly based on secondary sources (quoting basically other OECD reports, media and blogs) even though, these terms are not only common in restricted research circles, but in blogs and blawgs sustained by professional librarians, archivists and techno-lawyers. E.g. Robert Richards has kept his blog updated about resources, technical innovations and legal innovation systems (Richards, 2020).

A quick look at the terminology is also revealing. Surveillance is not mentioned in the report; privacy is mentioned only once, and data protection two times. These are real problems and hot topics regarding citizens. Government as a Platform was presented by O'Reilly[21] as "*a mechanism for collective action*" in the context of Obama Administration (Cass Sunstein was figuring out this kind of policy at the time, as Administrator of the White House Office of Information and Regulatory Affairs in 2010). However, it is clear by now that platform-based usage of RAC-type solutions can be monitored, controlled and surveilled without limitation, and this is what happened(see, Zuboff, 2019). The report does not address these issues (Draft, p. 70). Political participation and citizens' control of electoral processes, corruption etc. have been handled by independent platforms for years now (e.g. Ushaidi[22], GovRight[23] etc.). The same occurred with crisis and disaster management platforms (Poblet et al., 2017). Crowdsourcing, again, is not defined (nor mentioned) in the report, even if Openfisca is the result of a collaborative effort (logiciel open-source de micro-simulation du système socio-fiscal). Thus, it is not clear how the proposal of an open government is compatible with the "authoritative" version of rules and digital government.

There is an unsolved tension in the text between a top-down (authoritative) and a bottom-up (empowerment) approach. At the same time that the authors endorse a classical hierarchical normative legal order, they claim that:

> RaC can be seen as aligned with digital transformation, in that it envisages rules as digital instruments from the bottom-up, rather than as an add-on or an after-the-fact adjustment. If digital transformation is to succeed, then it will require the very basics of government to not just be digitised, but to be truly digital – thought of and built with digital technologies and mind-sets from the very beginning. RaC offers that possibility (Draft, p. 70).

Digitality and the implementation of AI and formal languages to ruling—the authors are correct—change the environment, relations and agency of governance. (In my opinion, 'Governance' would be a better notion to describe it). To avoid being trapped in a blind-alley, the report submitted in November 2019 to the EU Parliament On Good AI Governance by AI4People (Atomium Foundation)—Ugo Pagallo, Virginia Dignum, Robert Madelin, among many others— proposed a middle-out approach developed in 14 actions and a regulatory toolbox. We reflected the use of formal languages in a separated theoretical paper (Pagallo

---

[21]O'Reilly, T. (2010), "Government as a Platform", in Lathrop, D. and L. Ruma (eds.), Open Government, Chapter 2, https://www.oreilly.com/library/view/open-government/9781449381936/ch02.html.

[22]Ushaidi Crow Sourcing https://www.ushahidi.com/

[23]GovRight http://govright.org/

et al., 2019). It might be helpful if the authors considered a middle-out approach.

I deem RaC to be 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. Gartner has included one reference to machine-readable law (RaC), not in the 2020 hypercycle but understanding it as an example of digital twins for operational improvement (Finnerty, 2019). According to Wikipedia "A digital twin is a digital replica of a living or non-living physical entity. Digital twin refers to a digital replica of potential and actual physical assets, processes, people, places, systems and devices that can be used for various purposes". It is not yet clear, in this case, how the relation between "rules as code" and the existing legal instruments will be leveraged and combined. Finnerty points out that machine-readable legislation should be monitored and controlled, as it entails an organisational change at many levels of government (and this is a constitutional issue):

> Although it is expected that digital twins will develop graphical interfaces that support drag-and drop functionality, and that natural language programming or AI-application development (AI-AD; see "Innovation Insight for AI-Augmented Development") will eliminate the need to develop specific computer code, over time. However, comprehension of programming logic, data and modeling will still be essential, as will a natural language taxonomy for these government bodies. Whether machine-readable policy and legislation are developed by opposing political parties, other government branches or AI, similar comparative skills will be needed to evaluate them. Working from a single digital twin of the jurisdiction will require segmentation of users, code and datasets to support use by the different government bodies, such as a governor's policy office, and the various parties represented in a state legislature.

Finnerty also anticipates that "privacy, ethics and security concerns of citizens will challenge the use of digital twins of government".

> From a cultural perspective, there must be an understanding of the acceptance government leaders, political leaders and society have for machine-developed and -readable policy and legislation. There is a continuum from purely human-developed policy through to AI-developed policy. Adapting to the speed of change brought on by digital, while remaining thoughtful and determinant in developing policy and laws, will require governments to find the correct balancing point for their society. For most situations, this will mean a legislative and policy process that leverages augmented intelligence, but does not support full policy development and change through AI. To establish trust with constituents in this new approach, legislative and policy offices will need to evaluate the transparency of how laws and policies are developed and adapt to the needs of society.

I reproduce Gartner's graphic (Figure 3.1). It clearly shows the need of a broader discussion of government's (structural) digital policies.

The OECD Draft does not mention this digital twins policy which actually is what is at stake. It uncritically assumes its benefits:

> The codifications of laws and rules, in the sense of making them explicit and legible, is crucial for effective rules. It allows them to be known and shared, and encourages consistency in their application. Implicit rules, such as norms, are inherently harder to navigate and enforce. Explicit rules ensure there is, to some degree, a shared understanding and expectation of what is allowed or not. As digital transformation unfolds, many of these rules become more and more embedded in digital systems and structures. For instance, rather than knowing the details of tax law, many will simply rely on digital systems when completing their tax return, accepting that it is likely in compliance with the rules because the system said so. In this way, digital transformation can make explicit rules become implicit again – humans act in accordance with the rules embedded in digital infrastructure, even though the rules themselves may no longer be apparent or visible (Draft, p.24-25).
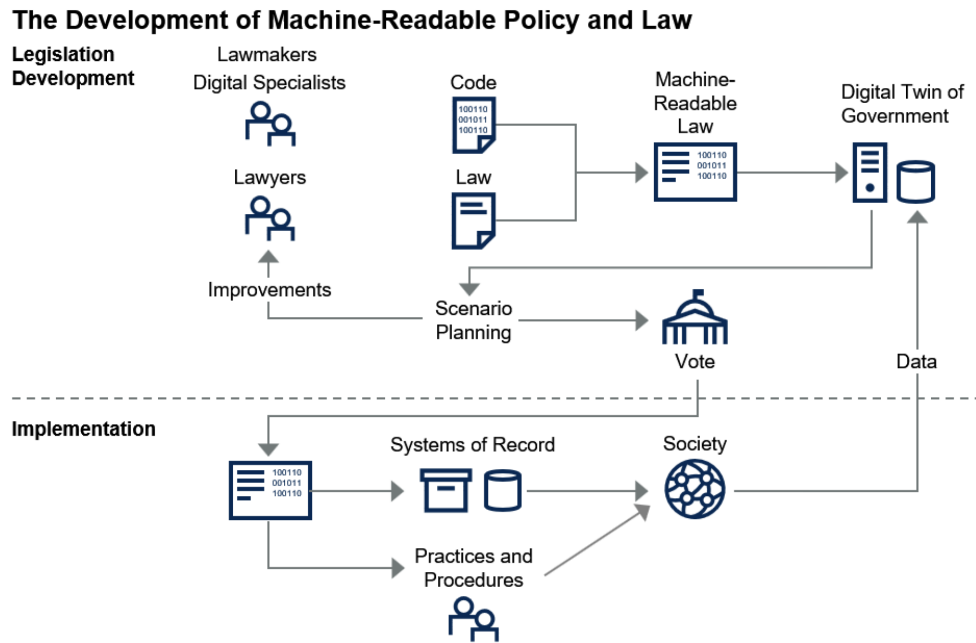
The Development of Machine-Readable Policy and Law

Figure 3.1: Gartner graphics (adopted from Finnerty, 2019)

This might be controversial. Norms are not implicit rules and are not "inherently harder to navigate and enforce". There is a controlled isomorphism, which has been discussed many times in the literature as a starting point for legal coding. And the iterative life cycling of the system is something that under all known methodologies must be monitored in all its development stages. Why should citizens accept that their tax return "is in compliance with the rules because the system said so"?

Sometimes the authors are unnecessarily inaccurate in their descriptions:

> RaC, in effect, forces rules to be explicit. It requires that rules are drafted in a manner that is explicit about the intent and interpretation of rules, as machines are as yet unable to engage in nuanced interpretation of ambiguity. RaC thus offers a structural driver for insisting that rules are drafted with greater clarity. In the absence of such a driver, with the rulemaking process being done by different people in different contexts, it is unlikely that rules will consistently be as clear as is desirable (ibid.).

Machines are able to cope with the problem of ambiguity, vagueness and rhetorical figures (i.e. metaphors, irony...) of language. This is not the issue: rules, a digital expression of norms, can respect the possibility of multiple interpretations; they can certainly suggest and create new meanings (see above, 3.3). What they cannot do is self-monitoring and evaluate their outcomes and impacts in a satisfactory way. And about the "desirability" of meaning: laws are created through and within a political process that is usually not peaceful. The ambiguity and vagueness of legal language is calculated, and it belongs to the classical legal toolbox (Bennett Moses, 2020). This is also the case of the so-called legal fictions (already theorised by Bentham and the analytical jurisprudence that followed).

I would have liked to see this document a bit more focused on social norms and social impact. The authors sparsely quote the Hammurabi Code, the Magna Carta and North's Institutions (1991). Some classical references in political science (I'm not referring to law, now) are missing and they are actually relevant to this subject, Mainly (i) Ullmann-Margalit, E. (1977). The emergence of norms. Oxford, Eng: Clarendon Press.; (ii) Elster, J. (2007, revised 2015) Explaining Social

Behavior: More Nuts and Bolts for the Social Sciences (Cambridge, 2007; revised ed. 2015); (iii) Conte, R., Andrighetto, G. and Campennl, M. eds., 2014. Minding norms: Mechanisms and dynamics of social order in agent societies. Oxford University Press; (iv) Ostrom, Elinor; Hess, Charlotte (2007). Understanding knowledge as a commons: from theory to practice. Cambridge, Massachusetts: MIT Press. (v) Aldewereld, H., Boissier, O., Dignum, V., Noriega, P. and Padget, J. eds., 2016. Social coordination frameworks for social technical systems (Vol. 30). Springer International Publishing.

(iii) and (v) can be deemed as results from the EU COST Program[24] SINTELNET[25] on collective intelligence. It is relevant because it delves into the structural coupling and de-coupling of artificial (cognitive) socio-technical systems internally and externally, within the social environment.

Following New Zealand report, the authors assume that the functional process of legal implementation occurs in a linear way:

> Yet, despite the apparent limitations of the theoretical model, the functional process required to move policy from development to implementation often accords with its basic tenets. Research on the 'policy intent' user journey from the NZ Government reveals that an often linear, sequential and siloed process underpins the movement from policy development to implementation. They write that: 'The current approach is relatively linear as Policy Development iterates with Ministerial Decision making and then moves to Legislative Development, before throwing the set of rules over the fence for operational implementation by Service Design and Delivery. If the policy is Operational then it skips the Legislative Development stages and goes straight into implementation.' (New Zealand Government, 2018).

They also assert its benefits:

> What RaC has the potential to offer is the ability to reduce the transaction costs of administering and complying with rules by reducing uncertainty and the need for costly analysis and interpretation, as well as potential contradictions between rulesets (Draft, p. 45).

I'm afraid this cannot be stated in this way. The "need for costly analysis and interpretation" is not explained, and there is no cost/benefit analysis offered to prove that RaC will effectively reduce transaction costs.

The Section 3.2 entitled ("*other preceding and related efforts*") of the Draft that mentions semantic languages could have reported the state-of-the-art. Well, these are not 'other' efforts but 'the' efforts, including 'legal'. The emergence of lawtech, regtech, fintech, insutech and suptech has occurred in the last ten years through legal web services. Computable models of the law and legal semantic web services were described and theorised many years ago, before that by Sartor, Prakken, Gordon, Governatori and many others (see, Casanovas et al. (2008), and Fernandez-Barrera et al. (2009)). The emergence of legal analytics (data mining, machine learning, etc.) did the rest (Nay, 2018). The application of logic to law is much older and goes back to the Middle Ages, Leibniz and (in the 20th c.) Ernst Mally and Layman E. Allen (Allen, 1957).

The authors write: "it has even been speculated that lawyers could soon be 'out of business'. While this is hyperbolic and, indeed, highly unlikely, the impact of technology in the legal domain is nonetheless significant." Instead of Forbes, the authors could have cited Phillip Susskind's books on this matter (over two decades) and the recent and precise Ashley (2017)'s account on AI & Law and legal analytics.

---

[24] https://www.cost.eu/

[25] https://ec.europa.eu/digital-single-market/en/news/sintelnet-european-network-social-intelligence

In relation to regulatory and legal compliance, the authors write:

> Following the Global Financial Crisis of 2007-08, for example, many governments assessed to what extent the regulatory frameworks and compliance measures governing the financial sector were sufficient. (Draft, p. 30)

> Coded rules, that is, rules in machine-consumable formats, already exist today. They are created by almost every enterprise that is required to comply with government regulations or legislation. They exist in the form of business rules, typically held in proprietary systems by each individual organisation. [...] . By centralising and making open the rules of government, and allowing third parties to consume these in an authoritative way, the need for interpretation (and therefore the risk of misinterpretation) is significantly reduced. It also assists in ensuring that changes are reflected in these systems in close to real time (Draft p. 57).

The Canadian experience description reads[26]:

> We can imagine a future where a new piece of regulation or legislation would be published not only in English and French but also in code. Doing this would make it easier for these rules to be consumed and interpreted by computers, allowing the creation of apps, software's and systems that have the coded rules built in. It is believed hat this would facilitate onboarding of digital services for government programs, improve compliance and reduce compliance costs. In addition, machine readable rules could allow policy makers and regulators to quickly and effectively model the outcome of proposed legislative or policy reforms using data and automated scenario testing as well as support automated or semi-automated administrative decision-making processes (for example, application forms and processing of applications).

Likewise, the OECD authors advocate for a "single provider of authoritative, machine-consumable rules" (Draft p. 76) to avoid interpretation and facilitate compliance:

> As conceived of here, RaC suggests that the actor best placed to provide a single and authoritative source of rules is the government. This represents more than the development of a new technical approach or technocratic 'fix' to an existing problem. It represents a potentially paradigmatic shift in the way the governments design, implement and provide rules. (Draft p. 77)

I already commented why this is not compatible with both formal (procedural) and substantive (rights) versions of the rule of law. But what matters now is explaining why companies and corporations started with the relatively new subject of compliance at the beginning of the century, (not after the economic crisis). Legal Compliance by Design is a term that was introduced to focus on the legality of the whole business process, mainly after the enactment of the Sarbanes-Oxley Act (2002), a US federal law that expanded and created new requirements for all public company boards and accounting firms after the Enron affair and other scandals. Business languages were developed in the nineties and the first decade of the 21st c. Many nuances were introduced—regulatory, normative, legal compliance; CbD, Compliance by Detection and by Default, etc. Many languages encompassing different assumptions were created for different purposes: Graph-based Business Process Modelling, Business Process Model and Notation (BPMN), Business Process Model and Notation – Query (BPMN-Q), Temporal Deontic Logic and Computational Tree Logic, Petri nets, etc. Several surveys on regulatory compliance have been already performed in the last ten years, including a meta-analysis of peer-reviewed systematic literature reviews on business process compliance (Akhigbe et al., 2015). One of the latest surveys are Hashmi et al. (2018b) and Hashmi et al. (2018a)—where the notion of (legal) 'Compliance through Design's is introduced. Several methodologies have been also developed for the legal domain: Regorous,

---

[26]https://wiki.gccollab.ca/CFR_Experimentation

Mercury, Eunomos, Nomos, Legal-URN and Legal Goal-Oriented Requirement Language, among others. Modelling legal compliance is challenging because of the complexity of legal system. There is no single, unique, avenue. What kind of compliance is the paper referring to? Definitions? Methodologies? The same for non-standard deontic logics or for the modelling of rights ('powers' in the Hohfeldian sense still constitute a problem).

The paper does not make the relevant distinctions. When the authors talk about CbD they also mention ethical compliance, and they certainly acknowledge the problem of turning legal norms into rules. They could also have paid attention to the different kinds of existing regulatory models in a globalised economy to counterbalance corporate models: responsible AI, transnational law, international courts, standards, protocols, human rights, international customary law, best practices, ISOs etc. But it seems that they resolve all jurisdictional and legal interpretative problems cutting the Gordian knot in one single blow: 'the actor best placed to provide a single and authoritative source of rules is the government'. Full stop. The consequence is that they get stuck into the narrow scope set by the notion of sovereignty in the nation-states and thus, into the classic and ancient notion of law.

About the French Government and their initiatives (as an example of "Rules as Code"):

> The French Government has a number of RaC-related initiatives. This includes the development of the open-source platform OpenFisca, LexImpact (which allows ex-ante policy modelling) and a number of France services based on coded rules (Mes Aides and Ma Boussoule).[27]

France has a strong tradition in AI and Law (Bourcier, 1995). These initiatives seem at first sight an intent to personalise and share the implementation of the law, facilitating the knowledge citizens' might have about their rights and duties according to the legal norms, and letting them share their experiences with the government and with their fellow citizens (Flückiger, 2019). I.e. a collaborative approach:

> Ma Boussole est là pour vous faire gagner du temps et de l'énergie en vous aidant à trouver des informations fiables et des aides personnalisées disponibles autour de vous.
> Ma Boussole vous donne accès à des témoignages d'aidants et d'experts pour vous aider au quotidien. Vous souhaitez partager votre expérience, témoignez maintenant!
> Ma Boussole est une plateforme collaborative qui évolue de jour en jour avec vous. Vous souhaitez nous faire connaître un acteur ou une solution qui n'est pas encore référencé, écrivez-nous un message.

I must still explore these development further, and I thank the authors for pointing them out. But, for what I know, this is a stick and carrot approach. The French state (still an administrative state, not to be confused with the Common Law and UK 'government') has responded quite harshly to LawTech services and 'private' initiatives performing legal analytics on sentences and the judiciary. It has been directly forbidden and penalised. Culprits can be punished with € 300.000 fine and 5 years jail.

> Les données d'identité des magistrats et des membres du greffe ne peuvent faire l'objet d'une réutilisation ayant pour objet ou pour effet d'évaluer, d'analyser, de comparer ou de prédire leurs pratiques professionnelles réelles ou supposées. La violation de cette interdiction est punie des peines prévues aux articles 226-18,226-24 et 226-31 du code pénal, sans préjudice des mesures et sanctions prévues par la loi n° 78-17 du 6 janvier 1978 relative à l'informatique, aux fichiers et aux libertés". Loi n° 2019-222 du 23 mars 2019 de programmation 2018-2022 et de réforme pour la justice.
> ELI: https://www.legifrance.gouv.fr/eli/loi/2019/3/23/JUST1806695L/jo/texte
> Alias: https://www.legifrance.gouv.fr/eli/loi/2019/3/23/2019-222/jo/texte

---

[27]https://www.maboussoleaidants.fr/accueil

One more observation: check the project on the Catala language based on default logic (Lawsky, 2017). It is not at INRIA. Project ID 15925 is listed as "deprecated" at `https://gitlab.inria.fr/verifisc/catala` and transferred to `https://github.com/CatalaLang/catala` (for some reason). The content of Box 6.1 could be more accurate. (Please clarify). I understand that it is developed by Dénis Mérigaux, stating that "The Catala language is special because its logical structure mimics the logical structure of the law." Asserting that the law has one (and only one - "the") logical structure, captured by a unique formal language is a strong epistemic assumption that should be justified.

4   Comments re Cracking the Code - Rulemaking for humans and machines

John Zeleznikow

Zeleznikow (2017) argues that algorithms have a role to play in supporting but not replacing the role of lawyers. He argues that while robots are unlikely to replace judges, automated tools are being developed to support legal decision making. In cases where litigants cannot afford the assistance of lawyers or choose to appear in court unrepresented, systems have been developed that can advise about the potential outcome of their dispute. This helps them have reasonable expectations and make acceptable arguments. Critics are concerned that the use of machine learning in the legal system will worsen biases against minorities or deepen the divide between those who can afford quality legal assistance and those who cannot. There is no doubt that algorithms will continue to perform existing biases against vulnerable groups, but this is because the algorithms are largely copying and amplifying the decision-making trends embedded in the legal system. In reality, there is already a class divide in legal access – those who can afford high quality legal professionals will always have an advantage. The development of intelligent support systems can partially redress this power imbalance by providing users with important legal advice that was previously unavailable to them.

Kannai et al. (2007) investigated how to best model judicial decision-making in order to use information technology to support enhanced legal decision-making in discretionary domains, They developed cognitive models of the exercise of discretion. They observed that discretionary decision-making can best be modeled using three independent axes: bounded and unbounded, defined and undefined, and binary and continuous decisions.

There has been extensive research in the development of decision support systems to model administrative justice, including the seminal work of Sergot et al. (1986) in interpreting the British Nationality Act of 1981. We believe that the use of code as rules to support legal decision making is most useful in administrative law domains. Discretion is closely associated with the concept of "open texture," a term first used by Waismann (1951) to assert that concepts are necessarily indeterminate. It is frequently used to describe the ambiguity or vagueness in the natural-language descriptions found in legal provisions or judgments.

Bench-Capon and Sergot (1988) define an open textured term as one whose extension or use cannot be determined in advance of its application. The distinct types of situations that are difficult to resolve because of the open textured nature of law are:

1. Classification difficulties,

2. Vague terms, and

3. Defeasible rules

Stranieri et al. (1999) note, their classification of legal tasks reflects domain experts' beliefs about four factors:

a. the extent to which the task contains ambiguous definitions;

b. the extent to which the predicates in the task are coarse);

c. the experts' jurisprudential perspectives;

d. the social and political environment.

Only tasks which are well defined (there are no open textured predicates) and bounded (decision makers have limited options) can be modelled using rules. Examples in this quadrant are drink driving in Victoria, Australia. and U.S.Sentencing Guidelines [USSG 18 USCS Appx.]. U.S. Sentencing Guidelines have actually been modeled by a (rule-based) computer system (Simon and Gaes 1989). Indeed, most jurisdictions dealing with driving infringements are rule-based and totally automated. This is also the case with the determination of social security benefits. In Australia this has led to the Robodebt problem. Sarder (2020) claims algorithmic decision-making has enormous potential to do good. From identifying priority areas for first response after an earthquake hits, to identifying those at risk of COVID-19 within minutes, their application has proven hugely beneficial. But things can go drastically wrong when decisions are trusted to algorithms without ensuring they adhere to established ethical norms. Two recent examples illustrate how government agencies are failing to automate fairness.

4.1   The algorithm doesn't match reality

This problem arises when a one-size-fits-all rule is implemented in a complex environment. The most recent devastating example is Australia's Centrelink "robodebt" debacle. In that case, welfare payments made on the basis of self-reported fortnightly income were cross-referenced against an estimated fortnightly income, taken as a simple average of annual earnings reported to the Australian Tax Office, and used to auto-generate debt notices without any further human scrutiny or explanation.

This assumption is at odds with how Australia's highly casualised workforce is actually paid. For example, a graphic designer who was unable to find work for nine months of the financial year but earned AUD 12,000 in the three months before June would have had an automated debt raised against her. This is despite no fraud having occurred, and this scenario constituting exactly the kind of hardship Centrelink is designed to address.

4.2   Inputs embed racism

Systemic racism has been repeated, more insidiously, in algorithmic processes. One example is COMPAS, a controversial "decision support" system designed to help parole boards in the United States decide which prisoners to release early, by providing a probability score of their likelihood of reoffending.

Rather than rely on a simple decision rule, the algorithm used a range of inputs, including demographic and survey information, to derive a score. The algorithm did not use race as an explicit variable, but it did embed systemic racism by using variables that were shaped by police and judicial biases on the ground.

Applicants were asked a range of questions about their interactions with the justice system, such as the age they first came in contact with police, and whether family or friends had previously

been incarcerated. This information was then used to derive their final "risk" score.

My argument is that except for very specific areas of administrative law it is unwise to develop legal rules as code.

5 Comments on Cracking the Code

Guido Governatori

The report fails to address an essential aspect related to the formal representation of norms: namely, it does not discuss the suitability of the language/formalism adopted for the representation. The field of Artificial Intelligence and Law has investigated such question for a long time. The idea of representing norms as in a formal language is not novel and traces back at least to the mid twenty Century. However, mostly, it remained a paper exercise till the 1980s with the so-called Imperial College approach proposing to use logic programming for the representation of norms as logical rules discussed in the seminal "The British Nationality Act as a Logic Program" paper (Sergot et al., 1986). Following the paper, there was a debate in the Artificial Intelligence and Law and Deontic Logic communities about different approaches/languages, mostly the descriptive approach advocated by the Imperial College school and the deontic Logic based approach. In 1991 Herrestad (Herrestad, 1991) showed that the descriptive approach is not adequate in general for the sound and conceptual representation of norms. In particular, the approach is not suitable when there are norms prescribing conditions in response to the violation of some obligations. While the descriptive approach is not suitable in general, there are applications where it is suitable, and it might be preferred to the more general approach. In general, languages and formalism for the representation of norms should be able to account for obligations, prohibitions and permissions (these are key concepts for the representation of norms), and at the same time, they are able to represent a large number of norms. In recent years, in the field of Business Process Compliance, several works proposed the use of languages based on Temporal Logic for the representation of norms (governing business processes). Temporal Logic provides whose semantics is similar to that of obligations, prohibitions, and permissions and the use of Temporal Logic for the verification of large scale industrial systems was the reasons why the fathers of model checking for temporal Logic have been awarded the Turing Award (the equivalent of the Nobel Prize in Computer Science). However, it has been shown (Governatori, 2015a; Governatori and Hashmi, 2015) that Temporal Logic is not able to correctly represent real life norms.

The examples above just demonstrate the subtleties required to formalise norms; in the AI and Law and Deontic Logic community we see over and over new proposals using technique proved to fail in the past. In the past few years, several approaches to Rules as Code have been proposed; however, most of them ignore the research done in the past on what are the real requirements for the appropriate formalisation of norms as computer code; in most cases not using the right approach failed the proposed projects. Some of the approaches discussed above would be suitable for specific applications; however, the designers for such applications should be aware of the limitations of the chosen approach and whether it meets the requirements for the specific applications. The report should discuss the requirements for the representation of norms, and the advantages/drawback of existing approaches/paradigms. Some starting points for this are in the following references.

Governatori et al. (2018) discussing rule-based declarative language versus procedural/imperative procedural (programming) languages in the context of blockchain. Thought the general discussion applies to the representation of norms in general. Gordon et al. (2009) discussing the general requirements for rule languages for the representation of norms. Finally, on the issue of interoperability, taking the fundamental issue of computation, something is computable by a computer if it is computable by a Turing machine. This means that two languages that are Turing complete can compute exactly the same things. In other terms, given the same input as a set of norms and facts, two Turing complete languages should produce exactly the same output. Very likely, the two languages will provide different functionalities for the representation of the norms. So the issue is about a common representation of the norms that can be used as an interchange format. To address this problem standard languages have been created. Specifically, here I refer to Akoma-Ntoso (LegalDocML)[1] and LegalRuleML[2] (both OASIS Standards). LegalRuleML is meant to provide a logic/language-independent formalism for the representation of norms. This means that the same representation can be used as an interchange representation and be shared by different applications (using internally, different logics or programming languages). Each application translates the representation provided in LegalRuleML to the format used internally by the application. The report should discuss the emergence of such standards and their benefits (and disadvantages) (Tara Athan, 2015) and the standard document[3] itself.

---

[1] http://docs.oasis-open.org/legaldocml/akn-core/v1.0/akn-core-v1.0-part1-vocabulary.html

[2] http://docs.oasis-open.org/legalruleml/legalruleml-core-spec/v1.0/cs01/legalruleml-core-spec-v1.0-cs01.html

[3] LegalRuleML core specification version 1.0 M Palmirani, G Governatori, T Athan, H Boley, A Paschke, A Wyner. Available at http://docs.oasis-open.org/legalruleml/legalruleml-core-spec/v1.0/csprd02/legalruleml-core-spec-v1.0-csprd02.html

6 Submission on Cracking the Code: Rulemaking for Humans and Machines

Mustafa Hashmi and Ho-Pun Lam

General Observation

As mentioned in the report, *Rule as Code* (RaC) could enable citizens and businesses to re-use and innovate using public infrastructure and can stand to other benefits envisioned by *Government as a Platform* (GaaP) — a concept based on a digital foundation for government to share data, software and services as an efficient, effective and innovation model for government (Margetts and Naumann, 2017). The report had spent a substantial portion of contents on the role of the government and the latest development of RaC in different countries. However, it should be noted that the "code" created by the government is in fact a form of law and can be a potent force for social liberation or control (McCullagh, 2009). It is important that lawyers, programmers, government officials and other actors involved should pay close attention to how such "code" should be formulated, and what kinds of services they are intended to provide and in what form. In addition, it is equally important to state what the limitations will be and what ought *not* to be done, i.e., something that is morally wrong or illicit.

6.1 Limitations of RaC: An example

A typical scenario of the application of RaC is to support decision-making process. That is, companies can make use of the encoded rules to develop applications that assist lawyers to determine the results of cases in hand and give meaningful legal opinion to their clients; whereas government can make use of the rules to develop applications to enhance their quality of services delivered to their citizens.

RaC, in essence, allows organizations to automate some of their decision-making processes. It is particularly suitable for YES/NO and IF-THEN-ELSE types of decisions-making. No doubt that the governments and organisations can hugely benefit from such automation for their decision-making processes such as to decide eligibility on tax benefit or employment support etc. However, from the decision-making and service execution perspective, the question is on the usability and effectiveness of the auto-generated decisions. Because in many situations YES/NO type answers might be not enough but further information may be needed to render a *'proper'* and *'just'* decision. Consider for example, the Heinz dilemma:

> A woman was on her deathbed. There was one drug that the doctors thought might save her. It was a form of radium that a druggist in the same town had recently discovered. The drug was expensive to make, but the druggist was charging ten times what the drug cost him to produce. He paid $200 for the radium and charged $2,000 for a small dose of the drug. The sick woman's husband, Heinz, went to everyone he knew to borrow the money, but he could

only get together about $1,000 which is half of what it cost. He told the druggist that his wife was dying and asked him to sell it cheaper or let him pay later. But the druggist said, *"No, I discovered the drug and I'm going to make money from it."* So, Heinz got desperate and broke into the man's laboratory to steal the drug for his wife. Should Heinz have broken into the laboratory to steal the drug for his wife? Why or why not?

In the above example, the algorithm consuming the coded rules may generate the decision of the person being guilty of stealing which is correct with respect to the provision stated in the legislation. However, what is lacking here is that it may not be stood right on the grounds of moral development in accordance with the levels of judgment (Kohlberg, 1981). Another example related to this is the *Al-Kateb v Godwin*[1] case which stated that the freedom of a person should not be deprived under any situation and is protected by the constitution (of Australia), which overturned the result made in accordance with the legislation.

As Reggie (2010) noted that the outcomes of such judgment may violate the basic law of society, deny others their rights, endanger the health and lives of others, or may involve the attempt to exploit others for personal benefits. Hence to prevent such injustice some additional information may be needed - for example, information on the individual circumstances or contextual information which lead to the wrong action, or some moral or ethical considerations. Yukl (2006) argues that bringing the considerations of both the individual circumstances and the contextual information than either variable alone (or without them) can better explain the situations - thus it can help in making better judgments.

Hence, this raises a question of the applicability of results generated by this type of legal system, i.e., under which situations the results can be trusted (That is, the weakness of such systems). Besides, this also raises another question on how such information will be made available and from where when the coded rules are used for the decision-making process? Besides, defining, and modelling the ethical considerations such as person's values, stage of moral development, freedom of choice, and the use of ethical or unethical behaviour etc., into the decision-making algorithms is extremely challenging, if not impossible. There is considerable disagreement about the appropriate ways to define, explain and represent ethics and contextual information, and evaluate the judgments based on ethics (Heifetz, 1994; Yukl, 2006). It seems, in addition to describe what the system can do, it is also needed to include some discussions on what the system cannot be done, which has not yet been covered in the report.

6.2   Technology

The discussion on Chapter 6 is a bit confusing. As mentioned above, it is vital that the parties involved should decide the types of services they are intended to provide before progressing to the next stage of development, which include justification on the technologies that should be used. Below are some of the questions that need to be considered.

6.2.1   System Architecture

System architecture focuses on the architecture of the system being developed and how it is going to be distributed. That is, is the system going to be a standalone system or library running on the users' machine, or is it going to be run on the government's (creator's) side as a web-/micro-services? How is the system going to be extended and scaled? Is it required to communicate, or interoperate, with other applications or external services? If yes, then how? etc.

---

[1] Al-Kateb v Godwin [2004] High Court of Australia (HCA) 37 219 CLR 562; 78 ALJR 1099; 208 ALR 124; Case Number: A253/2003, Retrieved from `http://eresources.hcourt.gov.au/showCase/2004/HCA/37`
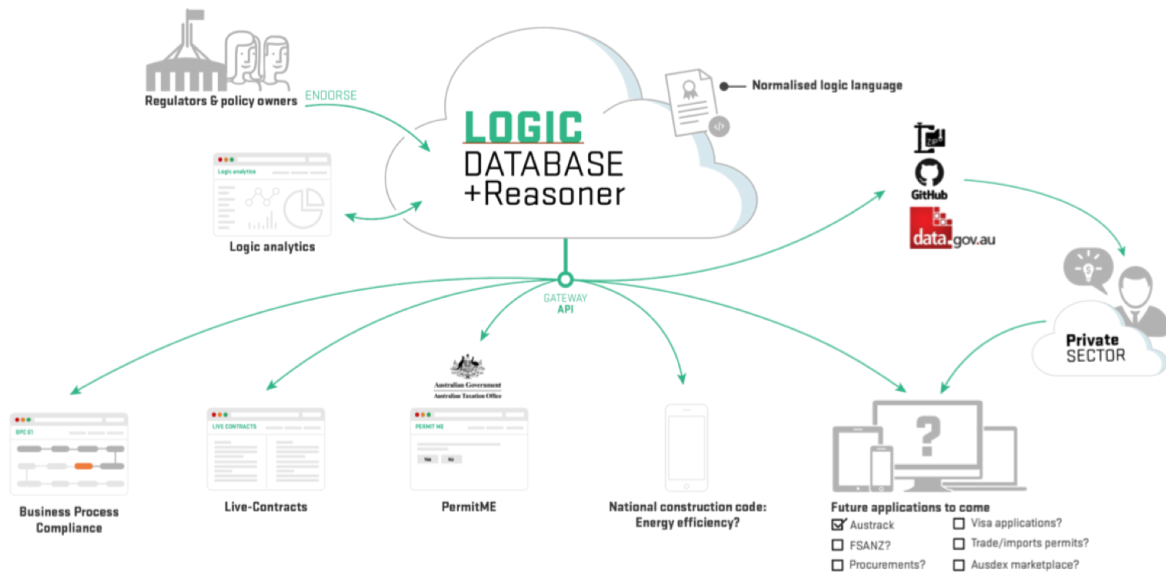
Figure 6.1: System Architecture of Regulation as a Platform at Data61, CSIRO (adopted from (CSIRO, 2018))

Consider Figure 6.1 below which depicts a high-level overview of the Regulation as a Platform (RaaP)[2] project architecture envisioned at Data61, CSIRO (CSIRO, 2018). As can be seen from the figure, legislations in RaaP will first be encoded using a (logical) formalism and stored in a database, and services can be performed, or provided to the clients, through invoking a set of APIs (application programming interface). Whereas systems like Clause and Xalgorithms are implemented as service portals that allows users to log-on to their systems and perform different tasks.

6.2.2 Modelling Language (or Formalism)

Modelling langauge (or formalism) concerns how legislation is going to be encoded (or modelled) and stored into the system. This is crucial as the expressiveness of the modelling language used will highly affect the usefulness and adaptiveness of the services being provided. It affects the effectiveness and correctness of intuitions that appear in the legislation being captured in the encoded rules, and the semantics that it represents. As a reference, a comprehensive list on the requirements on modelling legal rules and norms can be found in (Gordon et al., 2009), which includes the discussion on isomorphism, semantics, defeasibly, reification (such as jurisdiction and authority), contraposition, validity, legal procedures, normative effects, conflicts management, and traceability, of the rules.

In addition to this, as legislations change over time, in order to avoid any unnecessary duplication of the rules set, how to update the encoded rules becomes an important topic that needs to be investigated.

Besides, it is important to note that some legislations may also involve arithmetic and logic computations, as well as temporal information comparisons. In some cases, it may also need to connect/link to some external resources. Hence, how to encapsulate information like these into the modelling language is something that also needs to be considered when designing a

---

[2]https://data61.csiro.au/en/Our-Research/Our-Work/Future-Cities/Optimising-service-delivery/RaaP

language. Fortunately, some studies related to this, such as LegalDocML[3], has been published and is available in the literature.

Hashmi and Governatori (2018) argued that technical and structural complexity of the legal rules is an important factor for ascertaining the meanings of legal rules is far from being straightforward. With legal jargons, unintended and inconsistent interactions between different provisions of the same legislation (or across legislations) is a challenging task. To provide computationally efficient services such as compliance verification of legal rules, the modelling language should expressive enough to intuitively represent the legal rules. However, there is very limited evidence from literature (Hashmi et al., 2016) that such services may be intractable because of higher computational complexity of the formal language. Hence, for RaC to be more usable for providing efficient services, the chosen language needs to be expressive such that it is able to handle the set of large and highly complex legal rules.

However, it seems that these kinds of requirements have not been addressed clearly in the section in Section 6.2.1. The discussion of using imperative or declarative languages is too general and does not focus on the specific requirements that appear in legal knowledge (norms) representation. Languages such as Language for Legal Discourse (LLD) (McCarty, 1989), Legal Knowledge Interchange Format (LKIF) (Hoekstra et al., 2007), MetaLex (Boer et al., 2008), LegalRuleML[4] are some of the formalisms that are available in the literature.

6.2.3 Services support

Service support focuses on the types of services that the system would like to provide to the clients or intended users; while the services can be as simple as regulations retrieval, or advanced services which require inferring conclusions from the encoded rules under different scenarios (contexts).

It should be noted that, in addition to the three basic types of reasoning approaches (inductive, deductive and abductive), the use of case-based reasoning to retrieve similar cases that appear in common law (or precedent cases) is also an important area to study in the legal informatics domain, and some progress in this area has been achieved in the past few years (Grabmair, 2016).

It is correct that, at the moment, there is no single end-to-end solution available. This is due to the fact that, even though the same regulation is involved, the requirements of different applications (or different users) can be different, e.g., the requirements of conclusions being generated in query answering systems and recommendation or decision support systems can be totally different. However, irrespective of the services being supported, there are two main types of approaches in general. The approach used by RaaP and DataLex (AustLII, 2019) are commonly known as application agnostic approach as regulations are first encoded using a (logical) formalism and stored in a database, and transformations to other formalisms are made at a later time to accommodate the applications' need. It is flexible in a sense that regulation only needs to encode once and can be consumed by different applications. However, a trade-off for this is additional time may be required in performing the transformation.

Another type of approach is known as application dependent approach. As an example, in OpenFisca[5], the model that they developed is (mostly) restricted to the use of taxation computation (under different scenarios). That is, the applications developed using this approach are, in general, domain specific and may not be able to cater the needs of different

---

[3] https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=legaldocml
[4] http://docs.oasis-open.org/legalruleml/legalruleml-core-spec/v1.0/legalruleml-core-spec-v1.0.html
[5] https://openfisca.org/en/

34

applications. However, the applications developed can be very efficient as the code can be compiled in advance and stored in the system before use.

In summary, as RaC is intended to be provided as a public service, it is highly necessitates that the formalism employed should cater the needs of broader audiences (e.g., citizen and developers of different type of applications) in terms of *usability*, *adaptability*, and *accessibility*. Hence, an application agnostic approach seems more sensible and formalisms (languages), such as those mentioned in Section 6.2.2 above, can be good references to begin with.

In addition to the technical issues mentioned above, below are some of the questions that need to be considered but have not yet been discussed:

- Ownership of the encoded rules – who is the owner of the encoded rules? Is it going to be released publicly or proprietary (as trade secret)? And who is going to be accountable to these?

Bibliography

Okhaide Akhigbe, Daniel Amyot, and Gregory Richards. Information Technology Artifacts in the Regulatory Compliance of Business Processes: A Meta-Analysis. In Morad Benyoucef, Michael Weiss, and Hafedh Mili, editors, *E-Technologies*, pages 89–104. Springer International Publishing, 2015.

Layman E. Allen. Symbolic Logic: A Razor-Edged Tool for Drafting and Interpreting Legal Documents. *Yale Law Journal*, 66(6):833–79, 1957.

Giulia Andrighetto, Guido Governatori, Pablo Noriega, and Leendert W. N. van der Torre, editors. *Normative Multi-Agent Systems*, volume 4 of *Dagstuhl Follow-Ups*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013. URL `http://drops.dagstuhl.de/opus/portals/dfu/index.php?semnr=13003`.

Kevin D. Ashley. *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Cambridge University Press, 2017.

Rabeea Assy. Can the Law Speak Directly to its Subjects? The Limitation of Plain Language. *Journal of Law and Society*, 38(3):376–404, 2011. doi: 10.1111/j.1467-6478.2011.00549.x. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-6478.2011.00549.x`.

AustLII. Datalex: Legal inference systems, 2019. URL `http://austlii.community/foswiki/pub/DataLex/WebHome/DataLex_intro.pdf`.

Jeffrey Barnes. An AI or 'Rules as Code' Project: The Implications of Statutory Interpretation: A Short Comment. 2020.

Jerey Barnes. 'Shining Examples' - The Loophole 8, 2004.

Jerey Barnes. Sources of Doubt and the Quest for Legal Certainty. *Legisprudence*, 2(2):119–154, 2008. doi: 10.1080/17521467.2008.11424675.

Jerey Barnes. When 'Plain Language' Legislation is Ambiguous — Sources of Doubt and Lessons for the Plain Language Movement. *Melbourne University Law Review*, 34:671–707, 2010.

Jerey Barnes. *'The Plain Language Movement and Legislation: Does PlainLanguage Work?* PhD thesis, La Trobe Law School, La Trobe University, 2013.

Will Bateman. Algorithmic decision-making and legality: Public lawdimensions. *Australian Law Journal*, 94:520, 2020.

Trevor J. M. Bench-Capon and M. Sergot. Towards a rule-based representation of open texture in law. In Walter. C, editor, *Computer Power and Legal Language*, 1988.

Lyria Bennett Moses. *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence*, chapter Not a Single Singularity. Hart Publishing, UNSW Law Research Series No. 20-36, 2020. URL `https://ssrn.com/abstract=3646943`.

Francis A R Bennion. *Bennion on Statute Law*. Longman, 3rd edition, 1990.

Francis A R Bennion. Confusion over Plain Language Law. *The Commonwealth Lawyer*, 16: 63–68, 2007.

Tim Berners-Lee and Mark Fischetti. *Weaving the Web*. San Francisco: Harper., 1999.

Tim Berners-Lee, James Hendler, and Ora Lassila. The semantic web. *Scientic American*, 284 (5):34–43, 2001.

Guido Boella, Silvano Colombo Tosatto, Sepideh Ghanavati, Joris Hulstijn, Llio Humphreys, Robert Muthuri, André Rifaut, and Leendert van der Torre. Integrating legal-urn and eunomos: Towards a comprehensive compliance management solution. In Pompeu Casanovas, Ugo Pagallo, Monica Palmirani, and Giovanni Sartor, editors, *AI Approaches to the Complexity of Legal Systems*, pages 130–144, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg. ISBN 978-3-662-45960-7.

Alexander Boer, Radboud Winkels, and Fabio Vitali. MetaLex XML and the Legal Knowledge Interchange Format. In Casanovas et al. (2008), page 21–41. ISBN 9783540855682. doi: 10.1007/978-3-540-85569-9_2.

Danièle Bourcier. *La Décision artificielle: Le Droit, la machine et l'humain*. Presses Universitaires de France, Paris, France, 1995.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners, 2020. 020. Language models are few-shotlearners". arXiv preprint arXiv:2005.14165.

Pompeu Casanovas, Giovanni Sartor, Núria Casellas, and Rossella Rubino. *Computable Models of the Law: Languages, Dialogues, Games, Ontologies*. Springer-Verlag, 2008. doi: 10.1007/978-3-540-85569-9.

Pompeu Casanovas, Juan Arraiza Irujo, Felipe Melero, Jorge González-Conejero, Gila Molcho, and Montse Cuadros. Fighting Organized Crime Through Open Source Intelligence: Regulatory Strategies of the CAPER Project. In Rinke Hoekstra, editor, *Proceedings of the 27th Annual Conference on Legal Knowledge and Information Systems*, JURIX 2014, pages 189–198, Krakow, Poland, 2014. doi: 10.3233/978-1-61499-468-8-189.

Pompeu Casanovas, Monica Palmirani, Silvio Peroni, Tom van Engers, and Fabio Vitali. Semantic Web for the Legal Domain: The next step. *Semantic Web*, 7(3):213–227, 2016. doi: 10.3233/SW-160224.

Pompeu Casanovas, Louis De Koker, Danuta Mendelson, and David Watts. Regulation of big data: Perspectives on strategy, policy, law and privacy. *Health and Technology*, 7(4):335–349, Dec 2017. doi: 10.1007/s12553-017-0190-6.

M. Colesky, J. Hoepman, and C. Hillen. A critical analysis of privacy design strategies. In *2016 IEEE Security and Privacy Workshops (SPW)*, pages 33–40, 2016. doi: 10.1109/SPW.2016.23.

Commonwealth of Australia. Senate community affairs references committee, design, scope, cost-benefit analysis, contracts awarded and implementation associated with the better management of the social welfare system initiative. Senate Committee

Report, 2017. URL `https://www.aph.gov.au/Parliamentary_Business/Committees/Senate/Community_Affairs/SocialWelfareSystem/Report`.

CSIRO. Digital legislation: Tech solutions for red tape roadblock. Technical report, 2018.

Karin Derkley. Automation: Computer Says No. *Law Institute Journal*, 94:13, 2020.

Meritexll Fernandez-Barrera, Norberto Nuno Gomes de Andrade, Primavera De Filippi, Viola De Azevedo Cunha Mario, Giovanni Sartor, and Pompeu Casanovas. Law and technology: Looking into the future - selected essays. Pistoia, European Press Academic Publishing, 2009, Legal information and communication technologies; 7 Retrieved from Cadmus, European University Institute Research Repository, Florence: EPAP., 2009. URL `http://hdl.handle.net/1814/13482`.

Bill Finnerty. Governments Are Developing a Unique Kind of Digital Twin. Research Report G00450448, Gartner, Inc., 2019. URL `https://www.gartner.com/en/documents/3976081/governments-are-developing-a-unique-kind-of-digital-twin`.

Alexandre Flückiger. L'impact de l'intelligence artificielle sur l'écriture des lois: du code de lois à la loi encodée. *LeGes*, 30(3), 2019. URL `https://archive-ouverte.unige.ch/unige:128096`.

Stephen Gageler. Common Law Statutes and Judicial Legislation: Statutory Interpretation as a Common Law Process. *Monash University Law Review*, 37(2):1–15, 2011. URL `https://www.austlii.edu.au/au/journals/MonashULawRw/2011/17.html`.

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, editors, *Rule Interchange and Applications*, pages 282–296, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. ISBN 978-3-642-04985-9.

Guido Governatori. Thou shalt is not you will. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, ICAIL '15, page 63–68, New York, NY, USA, 2015a. Association for Computing Machinery. ISBN 9781450335225. doi: 10.1145/2746090.2746105. URL `https://doi.org/10.1145/2746090.2746105`.

Guido Governatori. The regorous approach to process compliance. In *8th International Workshop on Evolutionary Business Processes*, pages 33–40, Adelaide, Brisbane, September 2015b. IEEE.

Guido Governatori and Mustafa Hashmi. No Time for Compliance. In *2015 IEEE 19th International Enterprise Distributed Object Computing Conference*, pages 9–18, 2015. doi: 10.1109/EDOC.2015.12.

Guido Governatori, Mustafa Hashmi, Ho-Pun Lam, Serena Villata, and Monica Palmirani. Semantic business process regulatory compliance checking using legalruleml. In Eva Blomqvist, Paolo Ciancarini, Francesco Poggi, and Fabio Vitali, editors, *Knowledge Engineering and Knowledge Management*, pages 746–761, Cham, 2016. Springer International Publishing. ISBN 978-3-319-49004-5.

Guido Governatori, Florian Idelberger, Zoran Milosevic, Regis Riveret, Giovanni Sartor, and Xiwei Xu. On legal contracts, imperative and declarative smart contracts, and blockchain systems. *Artificial Intelligence and Law*, 26(4):377–409, 2018. ISSN 1572-8382. doi: 10.1007/s10506-018-9223-3.

Matthias Grabmair. Modeling Purposive Legal Argumentation and Case Outcome Prediction using Argument Schemes in the Value Judgment Formalism. June 2016. URL `http://d-scholarship.pitt.edu/27608/`.

Mustafa Hashmi and Guido Governatori. Norms modeling constructs of business process compliance management frameworks: a conceptual evaluation. *Artificial Intelligence and Law*, 26(3):251–305, 2018. ISSN 1572-8382. doi: 10.1007/s10506-017-9215-8. URL `https://doi.org/10.1007/s10506-017-9215-8`.

Mustafa Hashmi, Guido Governatori, and Moe Thandar Wynn. Normative requirements for regulatory compliance: An abstract formal framework. *Information Systems Frontiers*, 18(3): 429–455, 2016. ISSN 1572-9419. doi: 10.1007/s10796-015-9558-1.

Mustafa Hashmi, Pompeu Casanovas, and Louis de Koker. Legal compliance through design: Preliminary results. In Víctor Rodríguez-Doncel, Pompeu Casanovas, Jorge González-Conejero, and Elena Montiel-Ponsoda, editors, *Proceedings of the 2nd Workshop on Technologies for Regulatory Compliance co-located with the 31st International Conference on Legal Knowledge and Information Systems (JURIX 2018), Groningen, The Netherlands, December 12, 2018*, volume 2309 of *CEUR Workshop Proceedings*, pages 59–72. CEUR-WS.org, 2018a. URL `http://ceur-ws.org/Vol-2309/06.pdf`.

Mustafa Hashmi, Guido Governatori, Ho-Pun Lam, and Moe Thandar Wynn. Are we done with business process compliance: state of the art and challenges ahead. *Knowledge and Information Systems*, 57(1):79–133, 2018b. ISSN 0219-3116. doi: 10.1007/s10115-017-1142-1.

Ronald A. Heifetz. *Leadership without easy answers.* Cambridge, MA: Belnap Press of HarvardUniversity Press, 1994.

Luke Henriques-Gomes. 'robodebt: government admits it will be forced to refund 550m under botchedscheme',, 2020. URL `https://www.theguardian.com/australia-news/2020/mar/27/robodebt-government-admits-it-will-be-forced-to-refund-550m-under-botched-scheme`. The Guardian, 27 March 2020 (Australia).

Henning Herrestad. Norms and Formalization. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, ICAIL '91, page 175–184, New York, NY, USA, 1991. Association for Computing Machinery. ISBN 089791399X. doi: 10.1145/112646.112667.

Rinke Hoekstra, Joost Breuker, Marcello Di Bello, and Alexander Boer. The lkif core ontology of basic legal concepts. In P Casanovas, Biasiotti M. A, Francesconi E, and Sagri M.T, editors, *Proceedings of the Workshop on Legal Ontologies and Articial Intelligence Tech-niques, LOAIT 2007*, 2007.

Colin Howard. The language of the law. *Victorian Bar News*, 86:21, 1993.

Ruth Kannai, Uri Schild, and John Zeleznikow. Modelling the evolution of legal discretion: An artificial intelligence approach. *Ratio Juris*, 20(4):530–558, 2007. doi: 10.1111/j.1467-9337.2007.00374.x.

Dimitris Karagiannis. A business process-based modelling extension for regulatory compliance. In *Multikonferenz Wirtschaftsinformatik, MKWI 2008, München, 26.2.2008 - 28.2.2008, Proceedings*, 2008. URL `http://ibis.in.tum.de/mkwi08/17_IT-Risikomanagement_-_IT-Projekte_und_IT-Compliance/07_Karagiannis.pdf`.

Patrick Keyzer. Open rights or secret risk assessments? new challenges for publiclaw in an age of articial intelligence and the law. In Pompeu Casanovas Marta Poblet and Enric Plaza, editors, *Proceedings of the Workshop on Linked Democracy: Articial Intelligence for Democratic Innovation co-located withthe 26th International Joint Conference on Articial Intelligence*, 2017. URL `http://ceur-ws.org/Vol-1897/`.

Lawrence Kohlberg. *Essays on moral development / Vol. 1, The philosophy of moral development : moral stages and the idea of justice.* San Francisco ; Cambridge : Harper & Row, 1981.

Bert-Jaap Koops and Ronald Leenes. Privacy regulation cannot be hardcoded. A critical comment on the 'privacy by design' provision in data-protection law. *International Review of Law, Computers & Technology*, 28(2):159–171, 2014. doi: 10.1080/13600869.2013.801589.

Sarah B Lawsky. A logic for statutes. *Florida Tax Review*, 2017. doi: 10.2139/ssrn.3088206.

D Neil MacCormick and Robert S Summers. *InterpretingStatutes: A Comparative Study.* Dartmouth, 1991.

Helen Margetts and Andre Naumann. Government as a Platform: What Can Estonia Show the World? Technical report, Oxford Internet Institute, University of Oxford, 2017. URL `https://www.politics.ox.ac.uk/publications/government-as-a-platform-what-can-estonia-show-the-world.html`.

Patricia Marin-Chozas, Montiel-Ponsoda Elena, and Víctor Rodriguez Doncel. Language resources as linked data for the legal domain. *Knowledge of the Law in the Big Data Age*, 317:170, 2019.

L. T. McCarty. A Language for Legal Discourse I. Basic Features. In *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, ICAIL '89, page 180–189, New York, NY, USA, 1989. Association for Computing Machinery. ISBN 0897913221. doi: 10.1145/74014.74037.

Declan McCullagh. What larry didn't get, May 2009. URL `https://www.cato-unbound.org/2009/05/04/declan-mccullagh/what-larry-didnt-get`. Retrieved 15 Semptern, 2020.

James Mohun and Alex Roberts. Cracking the code: Rulemaking for humans and machines, discussion draft (oecd, 2020) par 8.2. 2020.

Ansgar Mondorf and Maria A. Wimmer. Requirements for an architecture framework for pan-european e-government services. In Hans Jochen Scholl, Olivier Glassey, Marijn Janssen, Bram Klievink, Ida Lindgren, Peter Parycek, Efthimios Tambouris, Maria A. Wimmer, Tomasz Janowski, and Delfina Sá Soares, editors, *Electronic Government*, pages 135–150, Cham, 2016. Springer International Publishing. ISBN 978-3-319-44421-5.

Eamonn Moran. The relevance of statutory interpretation to drafting. In *Drafting for the 21st Century, Conference at Bond University Gold Coast, 6-8 February (Law Reform Commission of Victoria and Bond University School of Law)*, 1991.

Julián Moreno-Schneider, Georg Rehm, Elena Montiel-Ponsoda, Víctor Rodriguez-Doncel, Artem Revenko, Sotirios Karampatakis, Maria Khvalchik, Christian Sageder, Jorge Gracia, and Filippo Maganza. Orchestrating nlp services for the legal domain. In *Proceedings of the 12th LanguageResources and Evaluation Conference, ELRA*, 2020a.

Julián Moreno-Schneider, Georg Rehm, Elena Montiel-Ponsoda, Víctor Rodriguez-Doncel, Artem Revenko, Sotirios Karampatakis, Maria Khvalchik, Christian Sageder, Jorge Gracia, and Filippo Maganza. Orchestrating nlp services for the legal domain, 2020b.

John Nay. Natural language processing and machine learning for law and policy texts, April 2018. URL `https://ssrn.com/abstract=3438276orhttp://dx.doi.org/10.2139/ssrn.3438276`.

New Zealand Government. Better rules for government discovery report. Digital.Govt.NZ, Report, 2018. URL `https://www.digital.govt.nz/dmsdocument/95-better-rules-for-government-discovery-report/html`.

OECD. Algorithms and collusion: Competition policy in the digital age. Technical report, The Organisation for Economic Co-operation and Development, 2017. URL `www.oecd.org/competition/algorithms-collusion-competition-policy-in-the-digital-age.htm`.

Ugo Pagallo, Paola Aurucci, Pompeu Casanovas, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Burkhard Schafer, and Peggy Valcke. *On Good AI Governance: 14Priority Actions, a S.M.A.R.T. Model of Governance, and a Regulatory Toolbox*. AI4People, 2019. URL `https://www.eismd.eu/wp-content/uploads/2019/11/AI4Peoples-Report-on-Good-AI-Governance_compressed.pdf`.

Hilary Penfold. Legislative Drafting and Statutory Interpretation. *The Judicial Review*, 7:471, 2006.

Katharina Pistor and Chenggang Xu. Incomplete law. *NewYork University Journal of international Law and Politics*, 35(4):931, 2003.

Marta Poblet, Esteban Garciá-Cuesta, and Pompeu Casanovas. Crowdsourcing roles, methods and tools for data-intensive disaster management. *Information Systems Frontiers*, 20(6):1363–1379, 2017. ISSN 1572-9419. doi: 10.1007/s10796-017-9734-6.

Valentina Presutti, Enrico Daga, Aldo Gangemi, and Eva Blomqvist. Extreme design with content ontology design patterns. In *Proceedings of the 2009 International Conference on Ontology Patterns - Volume 516*, WOP'09, page 83–97, Aachen, DEU, 2009. CEUR-WS.org.

Peter Quiggin. *Key Issues in Judicial Review*, chapter Statutory Construction: How to Construct, and Construe, a Statute, page 78. (Federation Press,, 2014.

Amin Rabinia, Sepideh Ghanavati, Llio Humphreys, and Torsten Hahmann. A Methodology for Implementing the Formal Legal-GRL Framework: A Research Preview. In Nazim Madhavji, Liliana Pasquale, Alessio Ferrari, and Stefania Gnesi, editors, *Requirements Engineering: Foundation for Software Quality*, pages 124–131. Springer International Publishing, 2020. ISBN 978-3-030-44429-7.

Davidrajuh Reggie. Modeling Ethical Decisions. In Shkelzen Cakaj, editor, *Modeling Simulation and Optimization*, chapter 17. IntechOpen, Rijeka, 2010. doi: 10.5772/8956.

R. Richards. Legal information systems and legal informatics resources. blogs (selected)., 2020. URL `http://www.personal.psu.edu/rcr5122/Blogs.html`.

Victor Rodriguez-Doncel. Natural language processing and semantictools: A roadmap towards publishing law as data. Technical Report DC25008: DC3-3, 2019.

Victor Rodriguez-Doncel, Cristiana Santos, Pompeu Casanovas, Asunción Gómez-Pérez, and Gracia Jorge. A Linked Data Terminology for Copyright Based on Ontolex-Lemon. In Ugo Pagallo, Monica Palmirani, Pompeu Casanovas, Giovanni Sartor, and Serena Villata, editors, *Proceedings of the AICOL International Workshops 2015-2017*, pages 410–423, 2018. doi: 10.1007/978-3-030-00178-0_28.

Monika Sarder. From robodebt to racism: what can go wrong when governments let algorithms make the decisions, 2020. URL `https://theconversation.com/from-robodebt-to-racism-what-can-go-wrong-when-governments-let-algorithms-make-the-decisions-132594`. last viewed 14 Oct 2020.

Sanford Schane. Language and the law (continuum), 2006.

J. Marek Sergot, Fariba Sadri, R. A. Kowalski, Frank. Kriwaczek, Peter. Hammond, and H. T.

Cory. The British Nationality Act as a Logic Program. *Communications of the ACM*, 29(5): 370–386, May 1986. doi: 10.1145/5689.5920.

Andrew Stranieri, John Zeleznikow, Mark Gawler, and Bryn Lewis. A hybrid rule - neural approach for the automation of legal reasoning in the discretionary domain of family law in australia. *Artificial Intelligence and Law*, 7(2):153–183, 1999. ISSN 1572-8382. doi: 10.1023/A: 1008325826599. URL `https://doi.org/10.1023/A:1008325826599`.

Monica Palmirani Adrian Paschke AdamZ. Wyner: Tara Athan, Guido Governatori. Legalruleml: Design principles and foundations. reasoning web. pages 151–188, 2015.

William Twining and David Miers. *How to Do Things with Rules*. Cambridge Univeristy Press, 5th edition, 2010.

M. Van Opijnen and Santos C. On the concept of relevance in legal informationretrieval. *Articial Intelligence and Law*, 25(1):65–87., 2017.

Matthew Waddington. Machine-consumable legislation: A legislative drafter's perspective – human v artificial intelligence. *The Loophole - Journal of Commonwealth Association of Legislative Counsel*, 2019.

Friedrich Waismann. *Veriability: Logic and Language*. Cambridge: Blackwell, 1951.

Gary Yukl. *Leadership in Organizations*. Pearson / Prentice-Hall, 6th edition, 2006.

John Zeleznikow. Don't fear robo-justice. Algorithms could help more people access legal advice, October 2017. URL `https://theconversation.com/dont-fear-robo-justice-algorithms-could-help-more-people-access-legal-advice-85395`. last viewed 11August 2017.

Shoshana Zuboff. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Prole Books, NY, USA, 2019.