

A Typology of Automatically Processable Regulation

Authors: Clement Guitton*, Aurelia Tamò-Larrieux*, Simon Mayer*

* University of St. Gallen

Accepted for publication in *Law, Innovation, and Technology*, 14(2), 2022.

Abstract

The possibility of encoding regulation to make it processable automatically by computers has been gaining attention within the legal discipline. With it, an abundance of terms has emerged as much as an array of academic discussions providing different examples, raising different concerns, while, unfortunately, having different premises in mind. This makes contributions within the field of what we refer to as '*automatically processable regulation*' difficult to compare with each other and research results hard to transfer among different research projects and groups. To overcome this problem, we propose a typology that enables researchers to locate their research project within the domain of automatically processable regulation, understand what issues might arise depending on where within the typology a project falls, and determine the relationship between projects. The typology revolves around three dimensions: the *primary aim* of the project, the *potential for divergence* of interests amongst stakeholders, and the *degree of mediation by computers*.

Keywords: Techno-Regulation; Automatically Processable Regulation; Computable Law; LegalTech; Algorithmic Regulation; Personalized Law

1 Introduction

The role of software in regulating human behaviour has long been recognized. Within this domain, the sub-field of regulation that is encoded *for the purpose of being automatically processed by computers* is gaining prominence. In this article we elaborate on different concepts that have been referred to in the literature as *computer-encoded law*, *computable law*, *computational law*, *legal tech*, or *legal artificial intelligence*. Multiple governmental and private research projects have emerged pushing the boundaries of such legal technologies for a diverse range of applications, where examples stretch from the creation of *automatically processable regulation*¹ to the creation of experimental software systems for dispute resolution.

However, due to the inconsistent terminology used, and the imprecisions this necessarily leads to, it is often unclear what scholars refer to when documenting the implementation of such systems, and when analysing and critiquing resulting issues and challenges. This lack of clarity hinders the scientific process by rendering research in the field incomparable across groups and disabling the transfer of conclusions from one (successful but isolated) case study to other use cases in the field. It also undermines efficient scientific debate about issues of automatically processable regulation, as non-functional properties (e.g., the degree of transparency or adaptivity) of specific implementations (e.g., whether a manually configured expert system or a machine-learning algorithm is used) are conflated with general statements about (societal) issues that affect the domain as a whole.

¹ For a comprehensive definition of automatically processable regulation, see Section 2.

To give an example, a recent article² warns about ‘computational law’ posing a challenge to the protections afforded by the rule of law. While the raised points are valid and important, the article has a narrow focus on the use of ‘computational law’ *by judges* and does not acknowledge that similar systems could, in other contexts, *support* the rule of law, for instance by increasing the *accessibility of law for laypeople*. The lack of distinction between the *different potential usage* of automatically processable regulation, by *different entities*, using *different enabling technologies* with *differing non-functional properties*, and in *different contexts*, then, however undermines the argument made. A further critique has been that: ‘[The encoding of legislation] has to foresee all potential scenarios and develop sub-rules that hopefully cover all future interactions – it must be highly dynamic and adaptive to address and confront what cannot easily be foreseen by way of unambiguous rules’.³ Yet, this represents an entanglement of code-driven law with futuristic concepts of tailored laws and micro-directives where all sub-scenarios would automatically be generated and encoded⁴ and thereby poses an obstacle when trying to identify which are the current, genuine, issues of automatically processable law in light of the distinctions introduced above. Yet another often mentioned criticism is that automatically processable regulation conflates the legislative or normative aspect of regulation with its execution or enforcement.⁵ This statement, while correct in many instances, however misses nuances as in other context such a conflation does not necessarily occur and depends largely on the extent of agency the system leaves to humans.

Beyond these examples, there have been notable discussions of a ‘legal singularity’, i.e., the idea that advancement in computer science will lead to only one singular and predictable interpretation of law.⁶ Although such futuristic, and potentially fictitious, scenarios that automatically processable law might lead us to in the extreme should not guide the scholarly discourse, it has become common to read about scepticisms and criticisms of legal singularity.⁷ While this scepticism is valid, we would like to highlight the underlying technical assumptions of such criticism that rely on ambiguous definitions and capabilities of machine learning algorithms. Were such definitions and capabilities more carefully explored, this would enable to better distinguish between what is currently genuinely being implemented in the domain of automatically processable regulation and what currently remains marketing talk.

Considering these entanglements, this article aims to give readers an overview of attempts towards automatically processable regulation, their emphasis, the applied technical approaches, and the generated issues. To do so, this article is, aside from an introduction and conclusion, divided into four main sections. In Section 2, we retrace the legal and technical evolutions that led to automatically processable regulation and introduce our terminology.

² Emilie van den Hoven, ‘Hermeneutical injustice and the computational turn in law’ (2021) 1 *Journal of Cross-Disciplinary Research in Computational Law* 1.

³ Mireille Hildebrandt, ‘Code-driven Law: Freezing the Future and Scaling the Past’ in Smon Deakin and Christopher Markou (eds), *Critical Perspectives on Law and Artificial Intelligence* (Hart Publishing 2020) p.70.

⁴ A. J. Casey and A. Niblett, ‘A Framework for the New Personalization of Law’ (2019) 86 *The University of Chicago Law Review* 333.

⁵ Mireille Hildebrandt, ‘A philosophy of technology for computational law’ in David Mangan, Catherine Easton and Daithí Mac Síthigh (eds), *The Philosophical Foundations of Information Technology Law* (Oxford University Press Forthcoming).

⁶ B. Alarie, ‘The Path of Law: Toward Legal Singularity’ (2016) 66 *University of Toronto Law Journal* 443.

⁷ S. Deakin and C. Markou (eds), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence* (Hart Publishing 2020); C. Markou and S. Deakin, ‘Ex Machina Lex: The Limits of Legal Computability’ in S. Deakin and C. Markou (eds), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence* (Hart Publishing 2020); J. Morison, ‘Towards a Democratic Singularity? Algorithmic Governmentality, the Eradication of Politics – and the Possibility of Resistance’ in S. Deakin and C. Markou (eds), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence* (Hart Publishing 2020); L.B. Moses, ‘Not a Single Singularity’ in S. Deakin and C. Markou (eds), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence* (Hart Publishing 2020).

Upon this basis, we elaborate in Section 3 on our main contribution to the literature: a typology of automatically processable regulation. We describe the different dimensions in detail and justify why these dimensions are important to frame the discussion. In Section 4 we then classify ten real-world examples of automatically processable regulation within our typology; these examples were initially discovered in the literature or through policy briefs and in order to provide more substantial (background) information on how the real-world projects evolved over time, we conducted a series of interviews with relevant stakeholders. This endeavour provides an in-depth discussion of applications of automatically processable regulation, shows how the typology works in practice, and enables researchers in the field to grasp the differences across implemented projects. Lastly, in Section 5 we discuss important findings from our analysis by addressing key questions that have been pondered by pioneers in automatically processable regulation and legal scholars.

2 The many facets of automatically processable regulation

2.1 Technology as regulation

Before elaborating on our typology, it is important to situate the topic at hand in the current state of research, and to define key terminology. Research on automatically processable regulation is not new and can be situated within the broader field of techno-regulation. Under the term 'techno-regulation' different relationships between technology and its impact on human behaviour are discussed in the literature. In the context of our analysis, two categories must be differentiated: First, technologies that by themselves regulate human behaviour without having necessarily been consciously designed to regulate in a particular manner – we refer to this as *technology as regulation*; and second, technology that encodes regulation, what we refer to as *automatically processable regulation*.

Technology as regulation should be understood as a broad category, since all technology in some way constrains or enables certain human behaviour. Thus, both terms, *technology* and *regulation*,⁸ should be understood broadly. Technology does not only include software-enabled applications; a well-known example often mentioned in this context is the design of overpasses in New York that has been inhibiting busses from driving through, thereby constraining poorer people dependent on public transport from easily accessing certain parts of the city.⁹ As technology can enable and prevent certain interactions and behaviour, multiple research threads on affordances and anti-affordances of technology have emerged.¹⁰ The second category of *automatically processable regulation* implements

⁸ Various definitions have been suggested within the literature. Famously Hood, Rothstein, and Baldwin (2001) define regulation as 'some standard setting element, some means by which information about the operation of the system can be gathered, and some provision for modifying behaviour to bring back within the acceptable limits of the system's standards'. C. Hood, H. Rothstein and R. Baldwin, *The government of risk: Understanding risk regulation regimes* (Oxford University Press 2001) p.23.

Or, Julia Black (2002): 'Regulation is the sustained and focused attempt to alter the behaviour of others to standards or goals with the intention of producing a broadly identified outcome or outcomes, which may involve mechanisms of standard-setting, information gathering and behaviour-modification'. J Black and D Kingsford Smith, 'Australasian Journal of Legal Philosophy [Plus a reply by Dimity Kingsford Smith]' (2002) 27 *Australasian Journal of Legal Philosophy* 1.

⁹ Langdon Winner, 'Do Artifacts Have Politics?' (1980) 109 *Daedalus* 121; Thomas J. Campanella, 'Robert Moses and His Racist Parkway, Explained' *Bloomberg* (Jul 9).

¹⁰ An affordance is the 'relationship between the properties of an object and the capabilities of the agent that determine just how the object could possibly be used' while an anti-affordance is the prevention of such a relationship and the interaction it enables. The concept of affordance provides for an attractive framework to more holistically analyse the regulation of technology. An interesting intersection emerges, when technology de facto creates or prevents interactions that are of legal relevance (e.g., use of Digital Rights Management (DRM) technologies to protect one's copyrighted materials, or the use of 'anti-

legitimately established regulation, e.g., autonomous cars that adhere to local traffic laws or implementations of data protection by design and default.

With *automatically processable regulation*, we refer to pieces of *regulation* that are possibly interconnected semantically and are expressed in a form that makes them accessible to be *processed automatically*, typically by a computer system that executes a specified *algorithm*, i.e., a *finite* series of *unambiguous* and *uncontradictory* computer-implementable instructions to solve a specific set of computable problems.¹¹ As the term indicates, automatically processable regulation has an element of 'automation' meaning that even if the coding process is manually done, the result is an automatic processing of a regulation. Furthermore, automatically processable regulation is implemented *intentionally*, i.e., there is a will of encoding regulation for a specific purpose, typically related to efficiency or accessibility gains.

More and more attention is being given to automatically processable regulation with an increasing number of books and journal articles on the subject matter emerging. In the literature, terminology that includes *algorithmic regulation*,¹² *computational* or *computable law*,¹³ *code-driven law*,¹⁴ *legal technologies (legal tech)*,¹⁵ *legal artificial intelligence*¹⁶, and *technological management*¹⁷ is in use. Because of the myriad of terms used to refer to different but overlapping concepts, it is key to be precise and distinguish between various ways to implement automatically processable regulation. For instance, Cobbe distinguishes technologies that *supplement* legal practice (legal technologies) from ones that *substitute* some legal actors or *depreciate* and *replace* traditional legal methods (legal artificial intelligence).¹⁸ This is an important distinction but might not be inclusive enough. First, it does not differentiate between *who creates* and *who uses* automatically processable regulation. This distinction is important when debating the legitimacy of such technologies and because certain conflicts of interest can emerge. Second, it seems to reduce the aim of automatically processable regulation to the *efficiency* gains within the legal domain, leaving out the potential *accessibility* aims that have been driving forces behind developments towards computer-encoded regulation. And third, Cobbe sets a focus on automatically processable regulation that builds upon *machine-learning*, even though many of the proposed implemented technologies follow (often manually implemented) rule-based approaches (see Section 4).

Even though in this article we use the term 'automatically processable regulation', we acknowledge that many of the mentioned examples below and the cited literature look into automatically processable *law*. In this context, law refers to *positively enacted legislation by a democratically chosen legislator*. The reason why we talk about *regulation* in the broader sense is that much of the push towards automatically processing regulation is coming from

homeless spikes' to protect residential properties from unwanted third parties). See: D. Norman, *The Design of Everyday Things* (Basic Books 2013); J. J. Gibson, *The ecological approach to visual perception* (Houghton Mifflin 1979); R. Leenes, 'Framing techno-regulation: An exploration of state and non-state regulation by technology' (2011) 5 *Legisprudence* 143; C. Gavaghan, 'Lex Machina: Techno-Regulatory Mechanisms and Rules by Design' (2017) 15 *Otago L. Rev.* 123.

¹¹ Math Vault, 'The Definitive Glossary of Higher Mathematical Jargon' (2021) <<https://mathvault.ca/math-glossary/#algo>> accessed May 25.

¹² Notably as edited book volumes: C. Busch and A.D. Franceschi (eds), *Algorithmic Regulation and Personalized Law: A Handbook* (Beck/Hart Publishing 2020); K. Yeung and M. Lodge (eds), *Algorithmic regulation* (Oxford University Press 2019).

¹³ Deakin and Markou (eds), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence*.

¹⁴ M. Hildebrandt, 'Code-driven Law: Freezing the Future and Scaling the Past' in S. Deakin and C. Markou (eds), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence* (Hart Publishing 2020).

¹⁵ M. Hildebrandt, *Smart Technologies and the End(s) of Law* (Edward Elgar Publishing 2015); J. Cobbe, 'Legal Singularity and the Reflexivity of Law' in S. Deakin and C. Markou (eds), *Is Law Computable? Critical perspectives on law and artificial intelligence* (Hart Publishing 2020).

¹⁶ Cobbe, 'Legal Singularity and the Reflexivity of Law'.

¹⁷ R. Brownsword, *Law, Technology and Society: Re-Imagining the Regulatory Environment* (Routledge 2019).

¹⁸ Cobbe, 'Legal Singularity and the Reflexivity of Law'.

the industry. In particular, proposals to encode safety requirements have emerged which ensure that artificial agents, such as manufacturing robots, comply with the local safety guidelines of a company or of a regulatory body. These regulations were not necessarily enacted by the legislature but can also reflect internal or industry policies.¹⁹

2.2 Evolution of legal expert systems

Human beings translating law into code has a long history. Some attempts go back to the mid-seventeenth century and point to the German mathematician Leibniz who in turn quotes Hobbes on how ‘everything’ is a math representation.²⁰ Even if such encoding attempts were rather at the meta-physical level, the quest for creating laws that can directly be interpreted by software has been ongoing (at least) since the 1960s.²¹ Important pioneer legal literature on this subject comes for instance from Layman²² and later together with Engholm²³ who defined legal drafting as ‘a mode of expressing ideas in statutes, regulations, contracts, and other legal documents in such a way that the syntax that relates the constituent propositions is simplified and standardized’.²⁴

In terms of actual implementation, the earliest example we encountered was the encoding of the British National Act in 1986.²⁵ Back then, four limitations prominently emerged, one due to the programming language used at the time, namely Prolog, and three due to the statute itself. First, Prolog did not accept rules that stated a negative conclusion; it would only infer ‘not P if fail to show P’. Second, the statute used ‘default reasoning’ in the absence of information, as well as counterfactual conditionals (e.g. ‘would have become a British citizen if hadn’t died’). Third, it used open-textured terms (e.g. ‘being a good character’). Such limits are by far not exhaustive. And the fourth limitation is syntactic ambiguity. Against this, a suggestion has been to identify each ‘atomic’ proposition and formalise (‘normalise’) it with propositional logic.²⁶

¹⁹ See for example A Shafei, J Hodgesc and S Mayer, 'Ensuring Workplace Safety in Goal-based Industrial Manufacturing Systems' (14th International Conference on Semantic Systems).

²⁰ The full quote states: ‘Thomas Hobbes, everywhere a profound examiner of principles, rightly stated that everything done by our mind is a computation, by which is to be understood either the addition of a sum or the subtraction of a difference... So just as there are two primary signs of algebra and analytics, + and –, in the same way there are as it were two copulas, “is” and “is not”’. G. W. Leibniz, 'On the Art of Combination' in GHR Parkinson (ed), *Leibniz: Logical Papers* (Clarendon 1666 [1966]) p.3.

²¹ A. Flückiger, 'L'impact de l'intelligence artificielle sur l'écriture des lois: du code de lois à la loi encodée' (2019) 30 *LeGes* unpaginated.

²² A.E. Layman, 'Symbolic Logic: A Razor-Edged Tool for Drafting and Interpreting Legal Documents' (1957) 66 *Yale Law Journal* 833.

²³ A.E. Layman and C.R. Engholm, 'Normalized Legal Drafting and the Query Method' (1978) 29 *Journal of Legal Education* 380.

²⁴ *Ibid.*

²⁵ M.J. Sergot and others, 'The British Nationality Act as a Logic Program. Communications of the ACM' (1986) 29 *Communications of the ACM* 370.. An even earlier seminal publication on legal expert systems is from 1981 attempting to implement automatic legal reasoning around product liability (and making a distinction between heuristic rules implemented in forward vs. backward chaining mechanism). However, this had more to do, again, with legal reasoning than with automatically processable regulation. Donald A. Waterman and Mark A. Peterson, 'Models of Legal Decision Making: Research Design and Methods' (1981) Rand Corporation (The Institute for Civil Justice).

²⁶ Layman and Engholm, 'Normalized Legal Drafting and the Query Method'.. Although the authors did not implement a system to do so, other researchers later did implement formal solutions, which could, in one case, even identify errors in the implementation by the state of a law regarding the computation of family benefits see e.g.: Jason Patrick Morris, 'Spreadsheets for Legal Reasoning: The Continued Promise of Declarative Logic Programming in Law', University of Alberta 2020); Denis Merigoux, Nicolas Chataing and Jonathan Protzenko, 'Catala: A Programming Language for the Law' (Forthcoming [submitted on Mar 4, 2021]).

This research resulted in 'legal expert systems' that reason on (manually) modelled knowledge and can overcome some of these limitations. For example, the *Process Compliance Language* implemented by the company *Neota Logic* models language for norm compliance, which can then be used by expert systems. *Process Compliance Language* supports four requirements: 'reasoning with defeasible legal rules, isomorphic linking from the logical rules to regulatory sources, expressing the kinds of obligations that statutes and regulations employ, [and] temporal reasoning'.²⁷ Many other languages have sprung up too – remediating certain challenges of Prolog as well as challenged that the British National Act raised. One of such languages has been most recently *Catala*,²⁸ aimed at possibly replacing the many old languages implementing calculations in the French administration, from benefits to taxes.²⁹ This leads to considering the question: What is currently technically possible? It helps to structure our answer by distinguishing between two distinct spheres: the one looking at encoding statutes, and the one looking at case law.

In statutes, one of the first steps is to determine terms in need of a definition, and to locate where such definitions could be found, as such definitions do not need to be tested against 'encodability'. References are often not clearly indicated, but still, many have already successfully mapped the interdependencies of laws within a specific corpus (e.g. EU law, Italian civil law) onto graphs.³⁰ Once there is an ontology of a statute, researchers may extract and encode definitions. An example thereof comes from Winkels and Hoekstra who have focused on the *Dutch Tax and Customs Administration* and created a model of all the different concepts in the very law (automatically) in order to form a basis for an implementation of the encoded version of the law.³¹ Other examples, from Francesconi and his colleague, have succeeded in classifying laws into eleven categories and with better results.³²

Once these definitions have been derived from a statute, the logic per se needs to be tested and extracted. Bach and colleagues applied techniques from graph theory on the *Japanese National Pension Law*, by first recognising logical parts in the law and then grouping them according to their common logical structure.³³ Their and other authors' success rates in

²⁷ M. Hashmi, G. Governatori and M.T. Wynn, 'Normative Requirements for Business Process Compliance' (Australian Symposium on Service Research and Innovation).

²⁸ See <https://catala-lang.org/>

²⁹ D. Merigoux and L. Huttner, 'Catala: Moving Towards the Future of Legal Expert Systems' (2020) hal-02936606 HAL.

³⁰ To do so, Koniaris et. al. focused not only on explicit references but also on subnetworks of regulations and legal basis. By mapping out close to 250 thousand legal texts from the EUR-Lex database, resulting in close to 1 million edges connecting the texts, they found out that 45% of the references would not show up by focusing only on explicit citations (or a built-up citation network). They also, interestingly, devised between those references which changed the meaning of the original document and those which did not. M. Koniaris, I. Anagnostopoulos and Y. Vassiliou, 'Network analysis in the legal domain: a complex model for European Union legal sources' (2018) 6 *Journal of Complex Networks* 243.

³¹ They used open source standards (CEN MetaLex) and vocabularies (SKOS), and were able to distinguish out of 167 thousand words, over 22 thousand words relating to concepts, and further extract a definition. Their conclusion was mixed with only 5% false positives but only 42% recalls (number of relevant concepts extracted and defined). They wrote somewhat optimistically: 'The automatic recognition of concepts in legislative texts is feasible, but not perfect (yet)'. R. Winkels and R. Hoekstra, 'Automatic extraction of legal concepts and definitions' (2012) 250 *Frontiers in Artificial Intelligence and Applications*.

³² For instance, Francesconi and Passerini in 2007 and Francesconi in 2009 obtained accuracy of 92.64% for classifying text, and to extract functional information with recalls averaging above 74% across different classes. Winkels and Hoekstra, in comparison, achieved recall rate of only 42% across 2 categories. Enrico Francesconi and Andrea Passerini, 'Automatic classification of provisions in legislative texts. *Artificial Intelligence and Law*' (2007) 15 *Artificial Intelligence and Law*; E. Francesconi, *An approach to legal rules modelling and automatic learning* (IOS Press 2009).

³³ For the first part, they looked for specific bits within articles to break them down into topic, antecedent and consequence, before then mapping it out onto a logical structure of requisite and effectuation parts

automatically recognising logic clauses was rather high (ranging from 75%-100%).³⁴ Further attempts for statutes have come from Benz Müller and colleagues³⁵ who used a specific 'value ontology' that tries to map balancing³⁶ norms within a two-dimensional plane along the axes *equality-utility* and *freedom-security* (oppositions which are debatable but which aim at simplifying the problems to try out solutions).

Lastly, when considering automatically encoding statutes, many limitations currently exist. For instance, Pertierra and colleagues tried to use default logic to parse a portion of the *US Internal Revenue Code* – a set of legislation heavy on calculations and that should, a priori, lend itself to such automatisation.³⁷ Yet, frequently, the parser misrepresented the 'or' operation when it should have been an 'and'; they had to manually code several key modules used as a stepping-stone for others; and a pool of experts needed to check the many ambiguities that arose. One of the reasons for these limitations is that most rules are not straightforward, and notably that 'tax law in its textual representation consists of long, complex sentences that are difficult even for humans to understand'.³⁸ The researchers used a combination of natural language processing (NLP) toolboxes that were then common in the field (C&C/Boxer) when attempting to grasp the semantic meaning of the code. Wyner and Governatori used it similarly too, and already warned of two related aspects: The need to manually check that the output of the parser matches the logic and semantics of the input, and more importantly, the gap between NLP and natural language when it comes to defeasible logic (amongst others) is still important enough that it is an area mentioned for future work in their publication.³⁹ In other words, current NLP technologies, while apt for short predicates, still fall short of grasping the complexity of (legal) language. In the words of Cobbe, it remains unreliable and unpredictable, to an extent at least.⁴⁰

(of the type 'if then', using a 'multi-layer sequence learning', in effect a binary classifier); for the second part, they used 'integer linear programming', a common technique for 'solv[ing] many NLP problems such as semantic role labeling', as the authors wrote. Ngo Xuan Bach and others, 'A Two-Phase Framework for Learning Logical Structures of Paragraphs in Legal Articles' (2013) 12 ACM Transactions on Asian Language Information Processing 3:1.

Interestingly, Takano et al. (2010) also focused on the Japanese National Pension Law, but on a specific sub-problem within it, namely the insertions within parentheses which modify prior meaning. They were looking to classify the logical structure into five main types and achieved correct recognition of the logical structure with rates varying from 75% to 100%. Kenji Takano and others, *Semantic Analysis of Paragraphs Consisting of Multiple Sentences: Towards Development of a Logical Formulation System* (2010).

³⁴ See Takano and others, *Semantic Analysis of Paragraphs Consisting of Multiple Sentences: Towards Development of a Logical Formulation System*; J. Zhang and N.M. El-Gohary, 'Automated information transformation for automated regulatory compliance checking in construction' (2015) 29 Journal of Computing in Civil Engineering.

³⁵ Christoph Benz Müller, David Fuenmayor and Bertram Lomfeld, 'Encoding Legal Balancing: Automating an Abstract Ethico-Legal Value Ontology in Preference Logic' (Forthcoming).

³⁶ Balancing norms refer to provisions that inherently require a balancing of different interests or establish requirements of 'reasonableness'. These norms cannot be applied in a clear-cut fashion and leave room for interpretation. See also C. M. Rose (1998) who differentiates among crystal and mud provisions (with a focus on property law). C. M. Rose, 'Crystal and Mud in Property Law' (1988) 40 Stan. L. Rev. 577.

³⁷ M. Pertierra and others, 'Towards Formalizing Statute Law as Default Logic through Automatic Semantic Parsing' (ICAIL).

³⁸ Ibid.

³⁹ A. Wyner and G. Governatori, 'A study on translating regulatory rules from natural language to defeasible logic' (Proceedings of RuleML).

Other more recent research still confirmed this on-going gap: Holzenberger and colleagues (2020) focused again on rules from the Internal Revenue Code and compared machine-reading models with hand-constructed ones. They conclude that 'machine reading models exhibits [sic] low out-of-the-box performance on our questions', and do not think that performance could be improved by applying it to larger datasets. N. Holzenberger, A. Blair-Stanek and B.V. Durme, 'A Dataset for Statutory Reasoning in Tax Law Entailment and Question Answering' (NLLP @ KDD).

⁴⁰ Cobbe, 'Legal Singularity and the Reflexivity of Law'.

Limitations are also present when trying to encode case law, but they are of a different nature. First is the lack of given structure, meaning that algorithms need to understand whether sentences are part of an argument or of a conclusion. Mochales and Moens⁴¹ obtained high accuracy rates when attempting to automatically classify sentences in two different sets of corpora (the *Araucaria* database⁴², and legal texts of the *European Court of Human Rights*, together making up more than six thousand sentences) with 74% and 80% respectively for sentences as part of arguments or conclusions. Others have also similarly used non-legal texts (briefs) to automatically classify whether an argument strengthened or weakened a side's claim, especially with the goal of predicting case outcomes (they achieved an accuracy of 91%⁴³) in the context of trade-secret misappropriation.⁴⁴ Another limitation is the need for manual annotations (although the need for annotation is also valid for statutes). This can be, for instance, to manually mark which sentences correspond to legal rule sentences or to evidence-based finding sentences, or to distinguish between facts and expert witnesses' opinions. Such annotation is, however, costly.⁴⁵

This field of research has given rise to two specific domains of applications: *information and argument retrieval systems*, and *e-discovery*. The former is concerned with giving an answer to a legal question and can be used by both laypeople and legal professionals. The latter is concerned with returning the right documents given a (legal) query: Corporate litigations and government investigations include a document review process during which each party goes over the documents relevant to the case, and hands them over ('produce') to the opposing party. It is possible that they withhold information by error (due to the oft-poor working conditions), or unlawfully on purpose (for which they could face sanctions by the court). Such reviews often involve assessing the relevancy of millions of documents and are time- and resource-intensive.⁴⁶ For both, information/argument retrieval systems and e-discovery, there are many projects running across academia and the private sector. Noteworthy in terms of limitation is that both need to rely on manual annotation, hindering automatization.

⁴¹ Raquel Mochales and Marie-Francine Moens, 'Argumentation mining' (2011) 19 *Artificial Intelligence and Law* 1.

⁴² Part of the *Araucaria* sets are non-legal texts (newspaper clips about the cases) which have been annotated.

⁴³ K. D. Ashley and S. Brüninghaus, 'Automatically classifying case texts and predicting outcomes.' (2009) 17 *Artificial Intelligence and Law* 125.

⁴⁴ Note that the above studies on encoding case law all used machine learning to attempt to automatically process the cases. For completeness' sake, it should be mentioned that they built on previous work that modelled arguments of past cases, but without using machine learning, rather with a logic and rule-based approach; the goal of such previous research was also often different and rather towards oriented towards information retrieval of cases that shared similarity with the facts or situation. The above studies on encoding case law all used machine learning to attempt to automatically process the cases. For completeness' sake, it should be mentioned that they built on previous work that modelled arguments of past cases, but without using machine learning, rather with a logic and rule-based approach; the goal of such previous research was also often different and rather towards oriented towards information retrieval of cases that shared similarity with the facts or situation.

⁴⁵ An example of the costs involved is given by Hendrycks and colleagues where they annotated 9,283 pages for legal contract reviews, at an estimated cost of USD 2 million. Even in the case of traffic regulations – an area of law a priori with less balancing norms – experts are needed to bring their judgement on whether specific instances would violate the regulations or not, for instance in cases of overtaking when the regulation specifies it can be made when safe, or with a clear view. See D. Hendrycks and others, 'CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review' (Forthcoming) p.2; H. Bhuiyan and others, *Traffic Rules Encoding Using Defeasible Deontic Logic* (IOS Press 2020).

⁴⁶ C. Privault and others, 'A new tangible user interface for machine learning document review' (2010) 18 *Artificial Intelligence and Law* 459.f

3 Typology

The long history of research on the topic of automatically processable regulation has meant that various projects have slowly come out of the realm of academia and have started being implemented by several state institutions and private companies – giving rise to much criticism on the effect that it might have on law and on the rule of law. Using a bottom-up approach, we looked first at such concrete implemented cases that apply automatically processable regulation so that we could extract knowledge from this data. When analysing current and past real-life projects applying automatically processable regulation (the projects themselves are described in Section 4 in greater detail), the following **dimensions** stand out.

First, each project on automatically processable regulation follows specific **aims** (Section 3.1). These aims can be divided into two primary types of benefits that the projects are trying to achieve: *Accessibility* to the regulation and *efficiency* of the application of the regulation. Projects seeking to increase accessibility will generally generate certain efficiency gains too, at least to some extent. The emphasis here is hence on the *primary* type of benefits sought after. To identify a project as primarily seeking an efficiency aim, we propose the double-question: *Can the beneficiary be a layperson as opposed to a legal professional? And if it can, does the sponsor present the project as benefiting the layperson rather than the state or legal professionals?*⁴⁷ If any of these is 'no', we classify the project as seeking efficiency rather than accessibility.

Second, there are various *potentials* for the **divergence of interests** between the roles actors on the project embody, namely between those sponsoring the project or commissioning the project (henceforth referred to as *sponsors*), those who are implementing the automatically accessible regulation (henceforth referred to as *implementers*), those it would bring benefit to (henceforth referred to as *beneficiaries*), and the *users*.⁴⁸ On the dimension of divergence of interests, the scale goes from *no potential for divergence* (Class 0) to *potential for large conflict of interests* (Class 4). How the measure is applied is described below in more detail (Section 3.2).

Third, the technology supporting the automatically processable regulation varies (Section 3.3). For the sake of providing a concise overview of this broad field in the context of this article, we distinguish between **different degrees of mediation by a human**, with the spectrum stretching from *fully mediated by humans* (corresponding to a 0) to *not mediated by humans at all* (i.e., fully mediated by computers; corresponding to a 4) at either extreme of the spectrum.

Each of these dimensions matters individually but also taken together a complex picture of projects implementing automatically processable regulation emerges. We posit that the different issues that automatically processable regulation begets embed these dimensions, and, more crucially, that an analysis of societal issues benefits from distinguishing along these dimensions as they spawn vastly different consequences as well as highlight different tensions. In the following, we elaborate on the specificities of the classification along each dimension for projects in the domain of automatically processable regulation. Subsequently, Section 4 further elaborates on each example individually and provides context on their implementation, benefits, as well as misfits and current status of operation.

⁴⁷ In other words, at least officially, do the benefits to the layperson outweigh those to the state?

⁴⁸ This list of stakeholders closely mirrors the one from Greene with beneficiaries, those with decision authority, those directly responsible for the program development, and only replacing 'the people disadvantaged by the program' by users.

J.C. Greene, 'Stakeholders' in S. Mathison (ed), *Encyclopedia of evaluation* (Sage 2005) pp.397.. For a more comprehensive list of stakeholders, see L.C. Ballejos and J.M. Montagna, 'Method for stakeholder identification in interorganizational environments' (2008) 13 Requirement Engineering 281.

3.1 Primary scope and aim of the project

Whether automatically processable regulation is pursued *primarily* to make it more *efficient* to determine legal facts, or to make law more *accessible* for citizens, impacts society very differently.⁴⁹ In that sense we see a tension between projects that try to make law more understandable to laypeople and thereby facilitate making use of their legal rights, and projects that aim at increasing the efficiency of legal processes. The root of the tension lies in the value we embed into the concept of democracy. On the one hand, ultimately, the goal of making law more *accessible* should lead to clearer laws (by removing certain unwanted ambiguities) and to more engaged citizens, hence to a more participatory form of democracy. On the other hand, speed, associated with *efficiency*, is not a value of democracy per se. On the legislative side, protracted negotiations, deliberations in which parties look for consensus are signs of a vibrant democracy, as much as on the judiciary side, judges carefully weighing different arguments and making sure that they comprehensively understand all the facts are equally signs of functioning judiciaries.

We acknowledge that not all projects follow either one or the other aim. Some projects will sit in-between and thus the aims should not be seen as mutually exclusive categories but as a subset of each other in which accessibility is subordinated to efficiency. This is because, by making the law more accessible, projects inadvertently lead to generating efficiencies as laypeople may not necessarily have to consult with legal experts but would be able to query the law on their own. But the efficiency gain is only secondary to the goal of first-and-foremost making the law more accessible.

3.2 Potential for divergence of interests in the project

We assume that a new sponsor (e.g., a policymaker) will be interested in understanding what issues may arise when commissioning or implementing a project on automatically processable regulation, and the magnitude of these issues. As identified by past research, the following factors do show prominence in determining the success of projects involving public-private partnerships: political support, public/community support, transparent procurement, trust, clarity of roles and responsibilities among parties.⁵⁰ These correspond to our different stakeholder types (sponsors, implementers, beneficiaries, and users).

These factors reflect that most projects involve various government agencies, as well as, often, private actors. Within such a set-up, several forms of conflicts of interest can emerge which can impede the realisation of the project. For instance, private actors as implementers might be interested in cost-effectiveness while the sponsor or (public) beneficiaries might not necessarily.⁵¹ Within the public sector, certain policymakers might be interested in initiating a change to score points with their hierarchy or to extend their powers, while other public servants will have an interest in keeping things as they are, especially if they perceive them as working well. The examples we outline in Section 4 illustrate the potential for differences and divergences when interests do not align between those who should primarily be the beneficiary from the implementation of automatically processable regulation, those who sponsor it, those who implement it, and those who use it. This led us to consider the following foundation for this second dimension:

f₁: *The fewer actors there are between beneficiaries/sponsors/implementers/users, the smaller is the potential for divergence of interests.*

This does not mean that a higher number of stakeholder types will *automatically* lead to higher diverging interests; merely that the *potential* for such a divergence is higher. We note that

⁴⁹ When we refer to 'citizens', we mean as residents in the country, and not in reference to any nationality

⁵⁰ Ballejos and Montagna, 'Method for stakeholder identification in interorganizational environments'.

⁵¹ E. Smith, T. Umans and A. Thomasson, 'Stages of PPP and Principal-Agent Conflicts: The Swedish Water and Sewerage Sector' (2017) 41 Public Performance & Management Review 100.

within the context of *contests*, there is little empirical evidence for both the self-interest axiom (actors behave for their own self-interest instead of the one of the group), or for the group-size paradox (large groups are less successful in reaching a goal than smaller ones).⁵² However and very importantly, bringing about automatically processable regulations does not fit with the typical example of hunt-the-stag-or-defect-for-the-rabbit. This is no contest with a winner and loser. And within such a context of public-private transactions, f_1 has its importance.

Furthermore, within the context of the state, many studies from organisational theory have made it clear that it would be highly misguided to consider public institutions as monoliths;⁵³ power politics within state institutions and intra-ministries with agencies vying for more resources and influence, the complex interplay of career public officials and politicians, and notably the extent to which goals align, matter too.⁵⁴ Divergences following 'turf wars' are observed rather than deducted from raw numbers as in f_1 , and follow a methodology based on qualitative analysis, which should help refine f_1 .⁵⁵ In case a qualitative analysis does not return evidence of conflict, this would be where f_1 is the most powerful, as an absence of evidence for conflict of interests is not the same as evidence of absence of such conflict.⁵⁶ To ensure that the potential for divergence is hence not under-assessed, in this very case, f_1 will function as a floor.

This leads to considering our second foundation for this dimension:

f_2 : *The more goals are aligned across and within the different set of actors, the less potential for divergence of interest there will be.*

These two foundations taken together lead us to *taking the maximum class between the two returned values from f_1 and f_2 , to determine the class of divergence of interest.* The reasoning why we believe that this way to determine the divergence of interests is valid rests on the potential for projects to fail should one actor defect. Being conservative in measuring the extent for potential divergence by taking the maximum of the classes hence appears as appropriate.

⁵² M. Kolmar and H. Rommeswinkel, 'Group size and group success in conflicts' (2020) 55 *Social Choice and Welfare* 777; S.F. Shen, E. Akçay and D.R. Rubenstein, 'Group Size and Social Conflict in Complex Societies' (2014) 183 *The American Naturalist* 301.

⁵³ B. Bowornwathana and O. Poocharoen, 'Bureaucratic Politics and Administrative Reform: Why Politics Matters' (2010) 10 *Public Organization Review* 303.

⁵⁴ B.G. Peters, *The Politics of Bureaucracy: An Introduction to Comparative Public Administration* (Routledge 2018).

⁵⁵ Although coming from a different field, namely medicine, a useful framework to support the translation of qualitative analysis into a measure of the extent of conflict of interest (not only inter unit, but more broadly) can be established by assessing the following: '(1) the likelihood that professional decisions made under the relevant circumstances would be unduly influenced by a secondary interest and (2) the seriousness of the harm or wrong that could result from such influence'. The authors of the framework go on to state that 'in assessing the likelihood of undue influence, it is reasonable to assume that the greater that the value of the secondary interest is (e.g., the greater that the size of the financial gain is), the more probable is its influence'. As mentioned above, it does not have to be financial gain, but this would also apply to extending an office's influence or convincing a constituency. See Institute of Medicine (US) Committee on Conflict of Interest in Medical Research, Education, and Practice, 'Conflict of Interest in Medical Research, Education, and Practice' in B. Lo and M.J. Field (eds), *Conflict of Interest in Medical Research, Education, and Practice* (National Academies Press 2009).

⁵⁶ This is a point most prominently driven by N.N. Taleb, *The Black Swan* (Penguin Random House 2010).

Table 1 Classifying the potential divergence of interests

Whichever returns the higher class	f ₁ : Number of non-overlapping actors	f ₂ : Degree of observed divergence of aims across and within non-overlapping actors
Class 0	0	Close to non-existent
Class 1	1	Small
Class 2	2	Medium-small
Class 3	3	Medium-high
Class 4	4	High

This class system reflects observations in the different scenarios (described in Section 4) where medium-large (Class 3) or even large (Class 4) conflicts of interests do emerge. For instance, making law more accessible to citizens by creating a system in which social benefits can be requested over a platform (by querying the law on social benefits) has led to state administrators complaining that they would create more work for their already fully-engaged staff, hence resulting in a conflict between the sponsor, an innovation unit within a ministry, and the actual government body administering social benefits. Understanding and explicitly acknowledging the potential of such conflicts is important for future projects to learn from past experiences.

3.3 Degree of mediation by computers in the project

The last dimension in our proposed typology is tied to the non-functional properties of different technologies that underpin systems in the field of automatically processable regulation. When arguing about the (societal) ramifications of automatically processable law, we propose that it is not helpful for the discourse on automatically processable regulation to categorize algorithms simply as ‘imperative vs. data-driven’ or even as ‘artificially intelligent vs. traditional’: On the one hand, there are (simple) machine-learning systems that are fully mediated by humans (e.g., based on manually curated data); on the other hand, regarding the final objective of an algorithm, purely imperative programming can be problematic as well, perhaps because of misguided specifications, the sheer complexity of a project, or unconscious or conscious biases of the programmers. In this sense, we discourage the use of terms such as ‘artificial intelligence’ to describe systems for automatically processable regulation. Such terms are not only too broad to permit a discourse about a specific system, but they are also prone to large interpretation differences by different communities and individuals. Rather, we propose that the central non-functional property that should be considered regarding automatically processable regulation (and especially regarding its legitimacy) is *to what extent humans are disintermediated within this process*. That is, we emphasize the degree of automation in terms of the involvement of humans in guiding the specification and execution of a computer program.

Regarding programming paradigms, *imperative programming* denotes the formulation of sequences of statements where each statement directly changes the state of the program – it thus focuses on describing *how* a program operates, and other programming paradigms such as *structured* and *procedural* programming derive from this paradigm. Imperative programs are typically written by human programmers, and data that informs how the program should be constructed (i.e., which statements should be put for maximum performance of the

program, or for maximum security) is *explicitly mediated by the human programmer* (e.g., through their training, experience, and intuition, including conscious and unconscious biases). Computer programs may furthermore include elements that make decisions based on *curated or uncurated data sources*. Such programs are referred to as *data-driven*, and they (partially) act upon given data rather than only depending on the given sequences of imperative statements. Consequently, the logic of such programs can be changed by editing the content of the data structures as well as by editing statements in the code. A simple form of this includes programs that make use of look-up tables (which are separate from the program code) as part of their execution; more complex forms include programs that use rule bases. With systems that are *implicitly mediated by a human*, we refer to data-driven programs that make use of data sources that are *curated* in the sense that a human has actively compiled them or is actively managing them; such data sources are typically small because of the requirement for such manual curation.⁵⁷ In this context, machine learning then *typically* refers to data-driven programs that construct internal models based on samples (so-called training data); by applying these models to new inputs, such systems are, in many cases (e.g., in the field of computer vision), able to categorize new inputs or make predictions about inputs without having been explicitly programmed to do this, by generalizing from the supplied data. While machine-learning systems can be based fully or partially on manually curated data, this manual curation is often not desired or even infeasible due to the sheer size of the underlying data source, which leads to algorithms that exhibit a *very low degree of mediation by a human* – in the case of unsupervised learning for instance, the system is presented with input data and required to self-organize and find patterns in the data set without even being supplied with labels for these patterns.

We argue that the *degree of disintermediation of humans* in the implementation and execution process of a system is a highly relevant technology-oriented parameter when classifying projects in the domain of automatically processable regulation. While such disintermediation has a long history in computer science – e.g., regarding compilers that are explicitly constructed to be highly conservative⁵⁸ – it has recently been spotlighted due to the advent of machine-learning systems that use little-curated or even uncurated training data (e.g., a vast amount of training pictures) which leads to issues around what is referred to as ‘algorithmic bias’.⁵⁹ The high performance of these systems on some problems (e.g., image classification or NLP) is today tied to deep-learning approaches – ‘black boxes’.⁶⁰ More generally, we thus posit that the disintermediation of humans during the creation of a system leads to systems that are less interpretable, i.e., the disintermediation undermines the ability to explain the system’s behaviour correctly post hoc.⁶¹ While for a purely imperatively programmed system, there is, typically, an individual who will be able to explain why the system behaved in a specific way at all times, this is not guaranteed for systems that increasingly remove the human from the curation of the data and/or the code.

With respect to our proposed typology, we then require a metric that renders these observations comparable across projects. This is in general hard to do for outside observers since, similar to our arguments about the potential for divergence of interests, a large part of the information required to generate the metric is not known. We therefore propose to *estimate* the degree of disintermediation of humans by considering the complexity of the

⁵⁷ Note that conscious and unconscious biases can also be present in such systems. However, in contrast to systems that are explicitly mediated by humans, biases in this case might derive from the underlying data source in addition to humans who curate this data source and/or create the internal logic of (partially) data-driven programs.

⁵⁸ S. Triantafyllis and others, 'Compiler optimization-space exploration' (International Symposium on Code Generation and Optimization).

⁵⁹ Brent Daniel Mittelstadt and others, 'The ethics of algorithms: Mapping the debate' (2016) July-December Big Data & Society 1; Sara Hajian, Francesco Bonchi and Carlos Castillo, *Algorithmic Bias: From Discrimination Discovery to Fairness-aware Data Mining* (ACM 2016).

⁶⁰ Davide Castelvecchi, 'Can we open the black box of AI?' (2016) 538 Nature 20.

⁶¹ Zachary C. Lipton, 'The Mythos of Model Interpretability' (ICML Workshop on Human Interpretability in Machine Learning).

implementation of a specific solution, where we take into account *domain* factors, *code* factors, and *data* factors: *Domain factors* refer to the complexity of the targeted problem domain, and of the specific problem in that domain – in our context, this is the complexity of the law that is being encoded. *Code factors* refer to the complexity of the software that is used when automating the processing (e.g., the size of the underlying code base).⁶² Finally, *data factors* refer to the level of curation of underlying data (if any) that is used to train a system, which is linked to the size and complexity of such data. Summarizing, we argue that complexity that may come in the form of many regulations, many lines of code, a large number of production rules, a large training set, or a large number of weights in an artificial neural network contributes to making systems less interpretable and, hence, implies a disintermediation of humans from the deliberation process of the system.⁶³ Each of these three sub-factors is scored from 0-4, we take the Euclidian distance (square root of the sum of the squares), and we round it to the nearest 0.5, meaning that this metric goes from 0 to 7.

Figure 1 summarises the three dimensions of our typology and the cases that we considered (see Section 4), positioned according to our methodology.

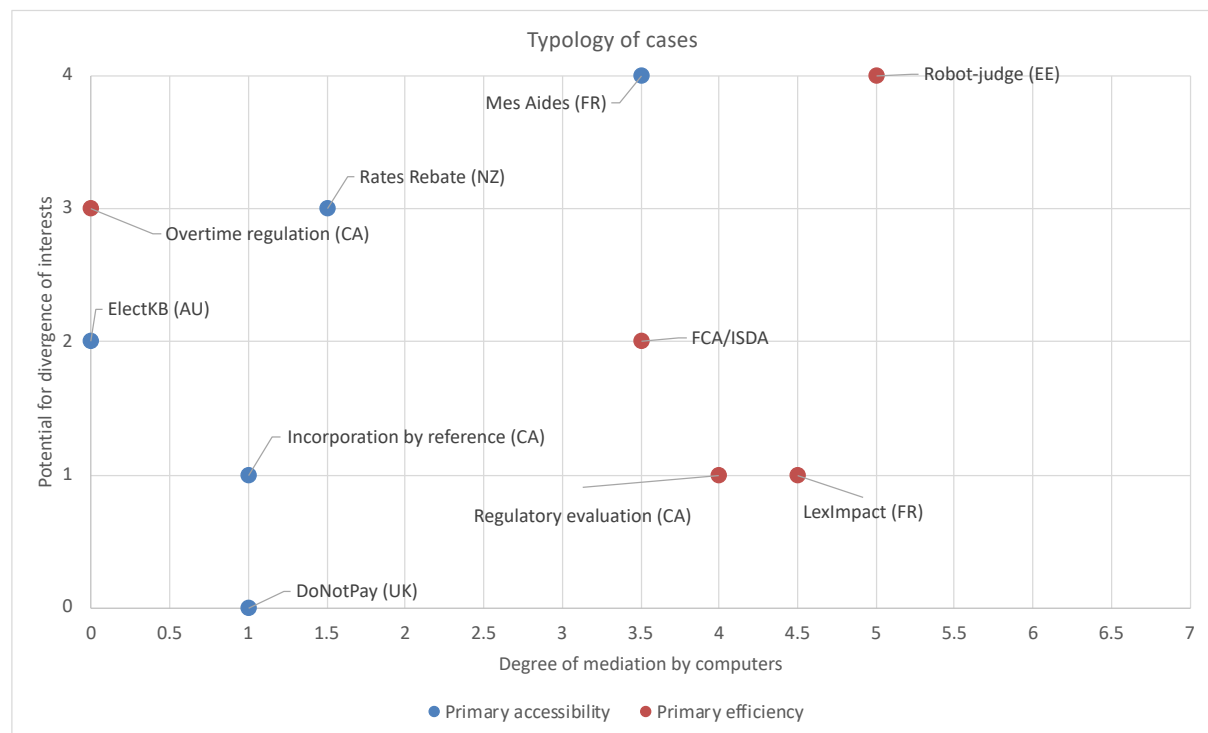


Figure 1: Typology of automatically accessible regulation

Overall, we posit that our proposed typology is key not only to address and take into account the individually mentioned dimensions (i.e., the underlying issues they trigger, the potential for conflicts that arise, and the interpretability and verifiability of decisions), but to better

⁶² See for software complexity for instance: Horst Zuse, *Software Complexity: Measures and Methods*, vol Walter de Gruyter (1990); E.J. Weyuker, 'Evaluating software complexity measures' (1988) 14 IEEE Transactions on Software Engineering; Joseph P. Kearney and others, 'Software complexity measurement' (1986) 29 Communications of the ACM 1044.

⁶³ A prerequisite to being interpretable in this sense is of course that the implementation of a decision-making system is accessible to inspection by stakeholders, and that stakeholders have the capability to understand the system in the first place (i.e., they need to be knowledgeable about software engineering or the training of artificial neural networks). However, our proposed typology takes a more detached viewpoint and attempts to evaluate the disintermediation of humans in principle. See Aurelia Tamò-Larrieux, Simon Mayer and Clement Guitton, 'How automation is changing the gatekeepers to law' *Medium* (Jan 26).

understand the interplay of dimensions. In fact, the mapping of cases helps to illustrate where the field and research on automatically processable regulation currently stands. It points to interesting findings such as the lack of *actually implemented* examples currently available with a high degree of mediation by computers, bar the robot-judge case from Estonia. Moreover, it helps identify areas where advocating for automatically accessible regulation may be more successful and create fewer issues, versus areas where lessons learned from other case studies can help us to avoid past mistakes.

4 Examples of automatically processable regulation

In the following we describe in greater detail ten real-world implementations from six countries that fall under our definition of automatically processable regulation. The intention of this section is not to provide a comprehensive review of all projects that are currently ongoing in this field – *although there are also only a few concretely implemented projects* –, but only of a selected subset that helps us showcase how our proposed typology applies to real-world implementations of automatically processable regulation. The projects below were well documented, or we were able to reach the involved parties and conduct interviews with them to fill in information gaps. The six interviews conducted were semi-structured, meaning that we followed a similar set of questions during all the interviews but left room for further topics that the interviewees pointed to.

To make the case studies comparable, the following elaborations will be structured according to the following guiding questions: We start with a brief description of the analysed project. Next, we discuss who the beneficiaries of the project are, and whether the sponsor presents the project as benefiting laypeople (to classify projects as accessibility-first or efficiency-first). Then, we describe the number of non-overlapping actors and the degree of observed divergence (to determine the potential for divergence of interests). Next, we discuss the technical implementation of the project with a focus on domain, code, and data factors (to determine the extent to which humans are disintermediated). And finally, we discuss other issues that we came across during our analysis.

4.1 Examples on accessibility and low mediation by computers

#1: DoNotPay (UK)	
Description	A private individual developed a 'robot-lawyer', which originally in essence prepares the right argument to advance in order to fight parking tickets. ⁶⁴
Accessibility vs. efficiency	The primary beneficiary is the public, and its sponsor (a private person) also presents it that way. We hence classify this project as aiming for <i>accessibility</i> first. Again, accessibility is defined with the benefits to <i>laypeople</i> that it brings. In this case, the efficiency gains for private citizens could also lead to more overload on the administration side due to them having to process more disputed parking tickets. This, in turn, could undermine the efficiency of the public service. Yet the focus on accessibility remains coherent with our definitions of who bears the gains (Section 3.1).
Potential for divergence	All sponsors, implementers, beneficiaries, and users are citizens. There are hence 0 non-overlapping actors, and no observed divergence of interest. Therefore, the degree for potential divergence is at 0.

⁶⁴ Z. Whittaker, 'Robot lawyer startup DoNotPay now lets you file FOIA requests' *Techcrunch* (Dec 10) <<https://techcrunch.com/2020/12/10/robot-lawyer-startup-donotpay-now-lets-you-file-foia-requests/>>.

Degree of mediation by computers	The law for parking tickets is mostly straightforward but variations heeding particularities per cities were introduced too. Thus, the domain factor is set at 1. <i>DoNotPay</i> was implemented manually without the use of past data. Thus, the code, and data factors are set at 0.
Other issues	Considering its success, this project has been extended to other areas, such as filing <i>Freedom of Information Access</i> requests and other legal questions. Reactions to these extensions have been mixed. ⁶⁵

#2: Incorporation by reference (CA)

Description	This project set out to analyse how many references to other documents there are within Canadian regulations (including to outside regulatory bodies) and how accessible these are to the general public. In the words of one of the project stakeholders: 'What does that mean for the rural Quebec farmer who is trying to be compliant, doing everything in their power to be compliant, but the standard we reference is only available in English [and behind a paywall, possibly]?' ⁶⁶
Accessibility vs. efficiency	As one of the public servants implementing it makes it clear, the public should be the beneficiary, and the sponsor (the School of Public Service) presents it that way. We hence classify this project as aiming for <i>accessibility</i> first.
Potential for divergence	The four roles are between the public service and citizens and there was no actual divergence observed; hence, we classify this project as a 1.
Degree of mediation by computers	The project looks for mentions of explicit references in different legal domains. While the algorithm must loop over all the legal texts and accurately find and interpret references, leading to a code factor of 1, the domain factor is 0 as the complexity of the system is low. Moreover, no past data is used as an input (data factor at 0).
Other issues	We encountered no other relevant issues.

#3: ElectKB (AU)

Description	This project aims at automatically checking the eligibility of individuals who wish to be elected to the Australian federal parliament. ⁶⁷
Accessibility vs. efficiency	Beneficiaries are citizens and the sponsors (universities) present it that way too. We hence classify this project as aiming for <i>accessibility</i> first.
Potential for divergence	Two universities developed it, without hints at conflicts of interest.

⁶⁵ Thomas Connelly, 'Computer science student, 19, says legal profession should be 'very scared' of his new 'robot lawyer'' *Legal Cheek* (Jan 15); Steph Wilkins, 'DoNotPay Is the Latest Legal Tech Darling, But Some Are Saying Do Not Click' *Above the law* (Oct 12) <<https://abovethelaw.com/legal-innovation-center/2018/10/12/donotpay-is-the-latest-legal-tech-darling-but-some-are-saying-do-not-click/>>.

⁶⁶ S. McNaughton, 'Innovate on Demand episode 7: Regulatory Artificial Intelligence' *Canada School of Public Service* (Jan 10) <<https://cspcs-efpc.gc.ca/podcasts/innovate7-eng.aspx>>.

⁶⁷ A. Mowbray, P. Chung and G. Greenleaf, 'Utilising AI in the legal assistance sector: Testing a role for legal information institutes' (2020) 38 *Computer Law & Security Review*.

Degree of mediation by computers	The election law that is being encoded is short and straightforward and a manual implementation was followed without relying on past data. Thus, all three factors are set at 0.
Other issues	We encountered no other relevant issues.

#4: Rates Rebate Act (NZ)	
Description	New Zealand's Service Innovation Lab implemented benefits calculators for two acts, the <i>Rates Rebate Act</i> and the <i>Holidays Act</i> . The calculator for the Rates Rebate is online and for everyone to use.
Accessibility vs. efficiency	Beneficiaries are citizens and the sponsor presents it that way too. We hence classify this project as aiming for <i>accessibility</i> first.
Potential for divergence	The centralised team in charge of looking to implement such 'system-wide' initiatives, the Service Innovation Lab, was dismantled after three years, as funding was cut, reflecting that the host ministry did not seem to value highly its conducted initiatives. Explanations for this have been linked to the divergence of views on the role of the public service. ⁶⁸ On the one hand, there are the proponents of seeing the public service as implementing minister's directives and as cutting costs. Within this view, there are little or no resources for innovation geared towards accessibility. On the other hand, people view the public sector rather as a way to build up tools for a society to strive as much as it can, and where, therefore, innovation has its place. ⁶⁹ There was a sense of regret that the view of the state as fulfilling a limited mandate won over the one of promoting access and understanding to laws. Following this dismantlement, we assess the potential for divergence as a 3.
Degree of mediation by computers	The law implemented focused on rebates, a twenty-page long law, with many definitions. A manual implementation was chosen without reliance on past data. Therefore, we set at domain and code factors at 1, and data at 0.
Other issues	<i>OpenFisca</i> is a tool to facilitate the encoding of automatically processable regulation. ⁷⁰ Many projects have leveraged it, including a few presented here. However, those using the platform leverage it only to create replicas of the actual implementation of the platforms taking the decision, not to replace them. In other words, implementers are re-coding what already exists, and citizens uses these re-implementations, while official administrative services keep using the old (legacy) versions. This means that there is a divide between those using it outside the officialdom of state organisations, and the actual system making a decision (the same is true for Mes Aides (FR) #6 below). There is a risk that citizens will not grasp the distinction between a simulator and the real

⁶⁸ Interview with stakeholder, March 23, 2021.

⁶⁹ Accounts described the impact of 'New Public Management' in the 80s and 90s which created silos within the public service as well as a lack of traceability of the law when implemented into computer code. These silos may have contributed to several scandals involving the public administration. One scandal in New Zealand for instance related to bad computer implementation of benefits, in its specific case, related to the under-, over-, and non-payment of 8,000 teachers, resulting in the Government owing NZD 12 million to teachers. See P. Andrews, 'New Public Management: The practical challenges, remedies, and alternatives' *The Mandarin* (Nov 21); L. Burr, 'Ministers' knowledge of Novopay problems revealed' *3News* (Feb 1).

⁷⁰ See <https://openfisca.org/en/>

	decision making system, especially that they will not understand the source for diverging output when it does not favour them.
--	--------------------------------------------------------------------------------------------------------------------------------

4.2 Example on efficiency and low mediation by computers

#5: Overtime regulation (CA)	
Description	This example from Canada centres on a piece of legislation, the <i>Motor Vehicle Operators Hours of Work Regulations</i> , which is rather short with only 20 sections running over two pages. The original scope of the project was to re-draft the regulation to modernise it, and while doing so to have a team working in parallel on the encoding of the regulation into programme code. Then, the aim was to post the resulting program code for consultation on a website and to implement an Application Programming Interface (API) and an application that would allow to test the impact of specific tweaks to the new regulations.
Accessibility vs. efficiency	The beneficiaries are mostly truck companies and regulators. While individual truckers can also monitor and test how their own working conditions are affected by it, it is assumed that they have little say on how their employers decide about their time allocation. Therefore, the focus does not rest on laypeople and the project falls under the primary aim of <i>efficiency</i> .
Potential for divergence	In this project, the original scope of work had to be curtailed: the part on updating the regulation by re-drafting it side by side with engineers had to be dropped. In Canada, in order to limit the potential for leaks which could be advantageous to a company learning in advance of changes to the regulatory framework, only individuals with a security clearance are allowed to be involved in the legislative drafting process. Such individuals are usually to be found on the policy side, not on the service side where public servants directly interact with the public and are faced with the on-the-ground implementation of regulations. The security clearance requirement means that, in practice, it is difficult to bring on-board both, firms to encode law (the technical expertise not being necessarily sufficiently developed in-house), and public servants on the <i>service</i> side. And hence, so far only legislators, lawyers, and policymakers have been involved in the process, maintaining the gap between the regulations and their implementation. In light of this, the sponsors had to drop the scope of trying to redraft a new piece of legislation while at the same time trying to turn it into an automatically processable regulation. Even without the redrafting process, the number of stakeholders involved is high, with Canada School of Public Service, policymakers, regulators, engineers, and truck companies. Hence, the mere number of actors make it correspond to a 3.
Degree of mediation by computers	The current (and non-updated) law is short and straightforward, and the implementation was carried out leveraging OpenFisca without using past data. Therefore, all the factors are set at 0.
Other issues	Another project from the same sponsor (Canada School of Public Service) did not go through and raises interesting issues: The individuals on the service side are fully allocated to their administrative duties and do not have legroom for innovative projects, even if these would afford them to be more efficient with their job. In this other project, a consortium of ministries in Canada intended to

	<p>work on encoding the rules for large commercial vessel registry.⁷¹ Key was that companies had difficulties in understanding the convoluted ways that vessels had to be categorised. The Department of Transport was willing to provide funding, but there was a misunderstanding. They thought that a different department, namely the policy and innovation branch of the Canadian School of Public Service, would work on an innovative solution and present it to them. The experience of the School though is that this does not work, and rather that the two teams needed to work 'linked to the hip as long as necessary', as an employee from the School describes it. After requesting that the Department of Transport allocates some time, the project was abandoned. This is a similar issue that we see raised in Mes Aides (see below, #6).</p>
--	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

4.3 Example on accessibility and medium mediation

#6: Mes Aides (FR)	
Description	This platform aims to offer an evaluation of the eligibility of individuals to 30 French benefit schemes.
Accessibility vs. efficiency	While the sponsors changed (it was first developed within an innovation unit of the state and is now maintained by a non-governmental organisation), the goal was always to benefit citizens and presented as such too. We hence classify this project as aiming primarily for <i>accessibility</i> .
Potential for divergence	Mes Aides exemplifies how observed conflict of interest can be high <i>within</i> the state. At the core has been the issue that Mes Aides did not have access to the exact code used by those taking the actual decisions; the exact code was closely kept as a secret by another ministry that was not ready to share it. ⁷² Mes Aide was trying to let citizens determine by themselves their eligibility, rather than to leave that only to the state administration, in the hope that more people who are eligible become aware of it and apply to schemes. The team behind Mes Aides therefore was required to re-encode the benefits law from scratch, using OpenFisca and a modern language (Python) which is however too difficult to integrate with older systems, and too slow to handle the millions of queries performed by the state officials. ⁷³ From the onset, it was clear that this would hence remain a simulation tool, and that the simulator would not cover all cases, crucially not when such cases would be too complex.

⁷¹ McNaughton, 'Innovate on Demand episode 7: Regulatory Artificial Intelligence'.

⁷² And this despite a law since 2016 (Oct 7) stating that the code for automated decision-making algorithms must be made public. Grounds to refuse publication involve a mixture of the following arguments: if someone were to discover a mistake in the code, a class action suit could cost dearly; it would add work to publish the code; or that it is too old to be usable by anyone nowadays. Arguments have been fuelling speculation around the real reason behind such a refusal is that the code is such in a poor state that it is embarrassing. The decision not to share the code is also probably counter-intuitive though, as by opening up the access to the algorithm to the public, it could allow for more researchers to explore and spot such mistakes early in the process, and in the end, could avoid the escalation of errors to a full-blown crisis decades later. To the credit of the French government, they have started doing so with selected university partners, as in the case of the computation for the income tax. See Inria, 'Avec Mlang, Inria participe à la modernisation du calcul de l'impôt sur le revenu' *Inria* (Mar 1 <<https://www.inria.fr/fr/mlang-modernisation-calcul-impot-revenu>>.

⁷³ Interview, Mar 19, 2021.

	<p>Opposing forces to the project put forward three general arguments:⁷⁴ It would turn the benefit-paying institutions into 'service providers'; it would overload them with queries; and it would increase fraud. On top of that came the worry that people who would obtain a positive answer on the simulator but a negative one from the real office would have difficulties understanding it, despite the online simulator stating as clearly as it could its nature. In the end, Mes Aides turned very popular: In 2019 alone, there were 2.3 million connections to the site, and 1,970 single visitors came on it more than a hundred times, hinting that professional social workers must have used the platform on behalf of citizens too, when trying to help them. And yet, despite this relative success and investments of €1.25 million across 5 years, the ministry abandoned the project in the year 2020, replacing it with a site merely providing textual information. However, a self-branded 'citizen community' reused the source code and copied it onto a new website, meaning that it is still accessible to the larger public. Regardless of how desirable the goal of making the law more accessible to the public is to ensure that as many people do understand it as they should, the French case for Mes Aides highlights a crucial point that is probably unavoidable: Power struggles within large bureaucracies and between ministries -- notably those acting as gatekeepers to the state secret and those seeking automated transparency. Due to these issues, we classify this project as a 4 on the potential for divergence domain.</p>
Degree of mediation by computers	<p>The corpus of law that was integrated was extensive, with each law embedding a different degree of complexity (domain factor: 3). This leads to a high complexity of the code (code factor: 2). No past data was used (data factor: 0).</p>
Other issues	<p>We encountered no other relevant issues.</p>

4.4 Examples on efficiency and medium mediation

#7: Regulatory Evaluation Platform (CA)	
Description	<p>This project aims at identifying conflicts between laws at the state, provincial, and city levels during periodic reviews.⁷⁵</p>
Accessibility vs. efficiency	<p>The beneficiaries are members of parliament and staff involved in the legislative process. Since users are professionals, we classify this project as aiming for efficiency first according to the double-no principle established above.</p>
Potential for divergence	<p>No potential for divergence was mentioned in our discussions.</p>
Degree of mediation by computers	<p>A vast scope of law is being covered (domain factor: 3), and the evaluation of whether conflicts in the law emerge is a complex undertaking (code factor: 3). No past data is being used (data factor: 0).</p>
Other issues	<p>We encountered no other relevant issues.</p>

⁷⁴ M. Alauzen, 'Splendeurs et misères d'une start-up d'Etat: Dispute dans la lutte contre le non-recours aux droits sociaux en France (2013-2020)' (2021) 225 Réseaux 121.

⁷⁵ McNaughton, 'Innovate on Demand episode 7: Regulatory Artificial Intelligence'.

#8: FCA/ISDA (UK/International)	
Description	In the UK, the Financial Conduct Authority (FCA) with the Bank of England has been looking into automating the implementation of financial regulatory requirements, so that in the end financial institutions can transfer to the authority several reporting data in a structured way. ⁷⁶ And similarly to what the FCA does, the Internal Swaps and Derivatives Association (ISDA) is looking to implement a similar process for the trading of interest rate and credit derivative products by seeking to promote a common standard 'in how derivatives trade events and processes are represented' as a result of regulatory changes. ⁷⁷
Accessibility vs. efficiency	The beneficiaries are regulators and private companies being regulated. We hence classify this project as aiming for efficiency first.
Potential for divergence	In general, such projects may not entail a high level of conflict of interest as both sides of the regulatory table see benefits to it: Regulators would be able to have standardised procedures and gain more visibility into regulated institutions; and regulated institutions would save on costs, on implementing changes in regulations, in less uncertainty (whether they could face fines because of their interpretation of the regulations, for instance), and in smoother communication with their regulators. This does not mean that the supervision is automated (which would arguably create a much higher potential for divergence), but merely the implementation of regulations. Hence, the number of involved stakeholders means that we set this dimension at 2.
Degree of mediation by computers	The financial regulations regarding capital and other regulatory requirements are complex, interwoven, and covering many different situations (domain factor: 3). Furthermore, although the code implementation occurs mostly manually, the complexity of the domain factor justifies a code factor of 2 (we assume that there are certain similarities in the computation of capital requirements across regulations, which justifies a 2 rather than a 3), with no use of past data (data factor: 0).
Other issues	We encountered no other relevant issues.

#9: LexImpact (FR)	
Description	The goal of this project is to 'help members of parliament to put a figure on their amendments and bills in the fields of tax law, public finance, and social affairs'. ⁷⁸ Some of the tools are publicly accessible (e.g., those relating to city budgets) while others are not (e.g., those related to the state's overall budget).
Accessibility vs. efficiency	The implementers, beneficiaries, and users are members of parliament, or staff of the legislative branch. We hence classify this project as aiming for <i>efficiency</i> first.
Potential for divergence	No potential for divergence was described.

⁷⁶ FCA, 'Digital Regulatory Reporting' *Financial Conduct Authority* (<www.fca.org.uk/innovation/regtech/digital-regulatory-reporting>).

⁷⁷ ISDA, 'What is the ISDA CDM' *International Swaps and Derivatives Association* (<www.isda.org/a/z8AEE/ISDA-CDM-Factsheet.pdf>).

⁷⁸ Assemblée Nationale, 'LexImpact' 2021) (<<https://leximpact.an.fr/>> accessed Mar 9).

Degree of mediation by computers	The laws for tax, public finances, and social affairs are complex (domain factor: 3). Furthermore, although the implementation is mostly manual, the process is highly complex (derived from the domain factor) justifying a code factor 3. And as the projects estimate both macro consequences of tax changes as well as impact on individual ‘average’ households, it has to use a variety of data and variables to be able to model this. We make the assumption that most of this data is uncurated as they would stem from general economic statistics, and therefore assign a data factor of 1.
Other issues	We encountered no other relevant issues.

#10: Robot-judge (EE)	
Description	Estonia gave the mandate to design a ‘robot-judge’ for disputes less than €6,400 in 2019. ⁷⁹ According to some, the label is misleading though, as it is rather automating the decision-making for uncontested claims. ⁸⁰
Accessibility vs. efficiency	This system can benefit citizens as they would have their case being decided faster and would still impart them with a better knowledge for the (application) of the law. But the way the main aim has been worded is rather to support the judiciary to work up on its long backlog of cases. Hence, we classify the project as aiming for <i>efficiency</i> first.
Potential for divergence	The mislabelling as ‘robot-judge’ instead of ‘automated decision-making for uncontested claims’ does carry further significance: It appears to be a reflection of a certain precipitation around the project, possibly driven by officials acting without considering key issues that it could potentially raise. In the words of two Estonian legal scholars: ‘the exciting image of AI overtook real events and created media interest in what turned out in fact not to be such a revolutionary change’. ⁸¹ After an initial pilot conducted with 65,000 trialled disputes, the project was discontinued based on the lack of funding available and possibly because of in-fighting between the Ministry of Justice, concerned with upholding the rule of the law, and the Ministry of Economy, with interests of cuts and efficiency at the forefront. But the project has been planned to reappear in 2021 following a prototype delivered to the Ministry of Justice, upon which more information may follow. Hence, we classify this as a 4 on the potential for divergence of interests dimension.
Degree of mediation by computers	The corpus of law that needs to be considered is complex, despite the low limit placed on claims, implying that the domain and code factors are a 3. Regarding the data factor, some assumptions are in order since only scant details about the actual implementation are known. We therefore assume that an engine would learn from the ruling of past cases. The project came to rule on roughly

⁷⁹ E. Niiler, ‘Can AI Be a Fair Judge in Court? Estonia Thinks So’ *Wired* (Mar 25).

⁸⁰ See K.N. Metcalf and T. Kerikmäe, ‘Machines are taking over -- are we ready?’ (2021) 33 *The Singapore Academy of Law Journal* 24. The way the original ‘robot-judge’ worked was that both parties need to agree that they delegate the judging to the robot. One of the critics has been thought, that at times, parties do not even agree on the facts, and that’s why it makes sense to focus first on ‘simple’ cases. The question as to whether the implementation of such robot-judges could be extended beyond such ‘simple’ cases remains to be seen. Cf. for the critic that parties do not agree on the facts Kevin D. Ashley, *Artificial Intelligence and Legal Analytics: new tools for law practice in the digital age* (Cambridge University Press 2017).p.51.

⁸¹ Metcalf and Kerikmäe, ‘Machines are taking over -- are we ready?’.

	75,000 cases ⁸² , meaning that the curation of this data set is large enough to allow a machine-learning system to generalize from it, and that this would require partial disintermediation of human. We hence also score the data factor component as a 3.
Other issues	The project raises many issues on the nature of law and its enforcement, which however at this stage go beyond the scope of this article.

⁸² Private communications with Estonian authorities, March 2021.

Table 2 Summary of the position of 10 projects according to the typology (* denotes the same unit across roles)

ID	Name	Country	Purpose	Sponsors	Implementers	Beneficiaries	Users	Aim	f1: # Non-overlapping actors	f2: degree of observed divergence	Potential for Divergence of Interests: $\max(f_1, f_2)$	Domain factor	Code factor	Data factor	Degree of mediation by computers (rounded to the nearest 0.5)
1	DoNotPay	UK	Avoid fines	Citizen	Citizen	Citizens	Citizens	A	0	Close to non-existent	0	1	0	0	1
2	Incorporation by Reference	Canada	Quantify open-access references	Public service*	Public service*	Citizens	Public service and citizens	A	1	Small	1	0	1	0	1
3	ElectKB	Australia	Check eligibility to parliament	Academia (2 unis)	Academia (2 unis)	Citizens	Citizens	A	2	Small	2	0	0	0	0
4	Rates Rebate Act	New Zealand	Check eligibility	Public service*	Public service*	Citizens	Citizens	A	1	Medium-high	3	1	1	0	1.5
5	Overtime regulation	Canada	For truck drivers	School of Public Innovation	Public-private partnership	Regulators and truck companies	Regulators and truck companies	E	3	Small	3	0	0	0	0
6	Mes Aides	France	Checking eligibility to benefits	Public service*	Public service*	Citizens	Citizens	A	1	High: Those in charge of innovation and those administering the benefits	4	3	2	0	3.5
7	Regulatory Evaluation Platform	Canada	Conflict between laws	Public service*	Public service*	MPs	MPs	E	1	Small	1	3	3	0	4
8	FCA/ISDA	UK/International	Financial regulations	Regulators	Regulators	Regulators and financial companies	Regulators and financial companies	E	2	Small	2	3	2	0	3.5
9	LexImpact	France	Quantify amendments	MPs	Unknown	MPs	MPs	E	1	Small	1	3	3	1	4.5
10	Robot-judge	Estonia	Disputes <€6,400	Ministry of Economy	Unknown	Ministry of Justice	Citizens	E	2-3	High: MoE concerned with cuts; MoJ with upholding the rule of law	4	3	3	3	5

5 Discussion

5.1 Does automatically processable regulation lead to more efficient regulation?

It is a fair assumption that private actors (e.g., lawyers, manufacturers) have an interest in efficiency gains. In line with this assumption are developments of e-discovery systems supported by advances in NLP,⁸³ as well as the foray from the autonomous vehicles industry to ensure that autonomous cars adhere to traffic regulation.⁸⁴ The examples above (Section 4) also indicate that the state apparatus has similarly an interest in projects increasing the efficiency of its activities relying on the application of law (e.g., judiciary and the Estonian *Robot-Judge* project), the determination of the impact of certain laws (e.g., *LexImpact* in France, *Overtime Regulation* in Canada), or the implementation of provisions for better accessibility to benefits (e.g., *Rates Rebate Act* in New Zealand).

The reasons why public and private actors are pushing for efficiency-driven implementations of automatically processable regulation are different *promises* attached to such projects. These promises include first, the ideal of reducing errors in the application of the law, second, the promise of reducing human biases by disintermediating them, and finally, the minimization of uncertainty and reduction of costs. We speak of promises because, as will be shown, it is highly debatable that such impacts are currently achieved.

The first promise is one of *reducing errors in the application of the law*. During the legislative process, it is common that certain ambiguities arise. While vagueness about terms – *semantic* ambiguity – may be introduced to smoothen diverging opinions,⁸⁵ there is however ‘no conceivable benefit’ from *syntactic* ambiguity⁸⁶ that arises due to a lack of clarity in the structure of sentences, and such mistakes also create a cost for society.⁸⁷ New Zealand reached a similar

⁸³ For instance, in one e-discovery system implementations, researchers tested their systems against a batch of 1,000 documents (‘a typical daily review amount’) from the Enron data and concluded with an overall precision of 88% (the system rather ‘over-assigned’, meaning that it tagged documents as relevant too often when in fact they were not). Privault and others, ‘A new tangible user interface for machine learning document review’.

⁸⁴ Note though, that it is unclear how much self-driving cars (e.g., of Tesla) can actually adhere at this point to traffic regulation. Cf. T. Krisher, ‘Tesla ‘full self-driving’ vehicles can’t drive themselves’ *AP* (Oct 22). Similarly, developments for autonomous ships, embedding an encoded version of maritime rules, exist, but the key core technical aspects are not public while the general process is. For an example of such processes, see for instance H. Ringbom and others, *Charting Regulatory Frameworks for Maritime Autonomous Surface Ship Testing, Pilots, and Commercial Deployments*, 2020.

⁸⁵ Some view this as a ‘feature’ of legal texts, see for instance: ‘The fact that written law affords such reconfiguration is not a bug but a feature of text-driven normativity, because it enables to calibrate and consolidate such mutual expectations in a way that is in line with past and future decisions – thus also weaving a fabric of legitimate mutual expectations that holds in the course of time’. Hildebrandt, ‘Code-driven Law: Freezing the Future and Scaling the Past’.

⁸⁶ Morris, ‘Spreadsheets for Legal Reasoning: The Continued Promise of Declarative Logic Programming in Law’.

⁸⁷ Consider the case of lorry drivers in the United States who brought a class suit in 2014 against their employer, Oakhurst Dairy. The local Maine law stated that employers were not to pay overtime in the case of: ‘the canning, processing, preserving, freezing, drying, marketing, storing, packing for shipment or distribution of: 1) agricultural produce; 2) meat and fish products; and 3) perishable foods’. The drivers

conclusion when attempting to update a piece of legislation while making it ready for encoding at the same time; they wrote that '[t]he process of developing rules statements identified gaps in the logic of legislation that had not been previously identified, which range from technical gaps that could be resolved by legal drafters through to gaps that require further analysis'.⁸⁸ Against this, a suggestion has been to identify each 'atomic' proposition, and formalise ('normalise') it with propositional logic.⁸⁹ Although the authors did not implement a system to do so, other researchers later did implement formal solutions⁹⁰ which could, in one case, even identify errors in the implementation by the state of a law regarding the computation of family benefits.⁹¹ As far as we are aware, we are yet to see a legislature *officially* implementing and using such a system though.

The second promise is one of *reducing human biases* by not involving humans. This argument has especially found resonance in the context of judges. Judges are (still) humans and, as such, are susceptible to the wide range of biases that affect humans. Such biases have for instance been demonstrated in a widely discussed study illustrating the impact of hunger on judge decisions.⁹² The main debate between legal formalists and realists has focused on whether judges can provide an impartial application of the law. In view of this uncertainty, the promise of projects like the *Robot-Judge* in Estonia is that automated judges would in fact not be susceptible to human biases.⁹³ The problem with this promise is that many algorithms are trained on biased datasets and have therefore reproduced old and current biases, for instance regarding racial bias.⁹⁴ We see the problem of biases also in another project in the Netherlands: In January 2021, the Dutch government resigned following wrongful accusations to 26,000 families of cheating the tax and benefit authorities. The accusations had led to families not receiving vital payment or worse, to being 'forced to repay tens of thousands of [E]uros, in some cases leading to unemployment, bankruptcies and divorces'.⁹⁵ Some of the accusations originated from an

argued that the law meant that the exemption is for 'packing for shipment' or 'packing for distribution', in the view of no commas separating the two, and that as a consequence, their employer owed them overtime incurred during distribution. They demanded USD 10 million -- an important sum for the 200-employee company which ended up settling with the employees for half that sum after the US Court of Appeals for the First Circuit confirmed the interpretation in favour of the plaintiffs. See D. Victor, 'Oxford Comma Dispute Is Settled as Maine Drivers Get \$5 Million' *The New York Times* (Feb 9).

⁸⁸ J. Stevenson, *Exploring Machine Consumable Accident Compensation Legislation: Lessons for a structural rewrite of the AC Act and opportunities to make it machine consumable*, 2019, p.31.

⁸⁹ Layman and Engholm, 'Normalized Legal Drafting and the Query Method'.

⁹⁰ Merigoux, Chataing and Protzenko, 'Catala: A Programming Language for the Law'.

⁹¹ Ibid.

⁹² An oft-quoted research looking at 1,112 cases found that Israeli judges examining parole cases in sequence were much more likely to deny it before a food break, and that the rate would return to its previous level after the break. The magnitude of the difference was massive: a case stood a 65% chance at the beginning of the series and would, so the authors, drop to near 0% when considered with a hungry stomach. The research was disputed on the ground that the ordering of cases was not random and several factors contributed to the downward trend in likelihood for acceptance (the case with higher chances would be those for prisoners with counsels and those would be lumped at the beginning of stacks), although the original researchers disputed such claims and countered that the data did not support it. Cf. K. Weinshall-Margel and J. Shapard, 'Overlooked factors in the analysis of parole decisions' (2011) 108 Proceedings of the National Academy of Sciences of the United States of America E833; S. Danziger, J. Levav and L. Avnaim-Pesso, 'Extraneous factors in judicial decisions' (2011) 108 Proceedings of the National Academy of Sciences of the United States of America 6889; S. Danziger, J. Levav and L. Avnaim-Pesso, 'Reply to Weinshall-Margel and Shapard: Extraneous factors in judicial decisions persist' (2011) 108 Proceedings of the National Academy of Sciences of the United States of America E834.

⁹³ V. Southerland, 'The Intersection of Race and Algorithmic Tools in the Criminal Legal System' (Forthcoming) *Maryland Law Review*.

⁹⁴ Ibid.

⁹⁵ J. Henley, 'Dutch government resigns over child benefits scandal' *The Guardian* (Jan 15).

algorithm wrongly flagging people for fraud, although as importantly, processes were set up in a discriminating way too, and the government failed to act upon reports from oversight bodies.⁹⁶

The third promise is one of *minimizing uncertainty and saving costs*. In particular for more rule-based processes, saving costs is at the forefront of many private companies and state authorities. States could render automatically processable laws like the Rebates Act in New Zealand or the Motor Vehicle Operators Hours of Work Regulations in Canada, and could provide them to other authorities and private companies. Thereby, companies and other institutions that are required to implement such laws could do so via digital interfaces that enable employees or citizens to have access to the legally guaranteed benefits they are allowed to obtain. However, also this last promise has not been achieved across the board. The *robodebt* project in Australia is a case in point. This project has attempted since 2016 to fully automatically recover social benefits that the state overpaid, a realisation made possible only once the different Australian ministries reconciled the income data reported to the social benefits department on the one side with the data from the tax department on the other side. This system recovered funds based on false information in 470,000 cases but amounting to an estimated AUD 1.5 billion in funds *wrongly* recovered out of AUD 2.1 billion collected -- and which will be paid back to the affected individuals.⁹⁷

To conclude, in light of these hurdles, a complex picture emerges between what projects aim and promise to achieve and what in the end the result is; automatically processable regulation does not lead necessarily to more efficient regulation, even when aiming at it.

5.2 Does automatically processable regulation lead to better access to regulation?

While in an ideal world (and following the rule of law), anyone should have access to and understand the law, the reality is that except for those very few who have had a specific legal training, the vast majority of citizens would face large hurdle to know where to find the relevant statutes and case law, and even more problematically, to know how to interpret them correctly in the right context, while navigating the winding legal processes. That is one of the key reasons why legal professionals have a business case and can charge high prices: Because they have trained and have experience navigating in this environment where few have.

Making the law more *accessible* to citizens brings us closer to the ideal concept that everyone should know the law. However, this ideal is by far not set in stone and the evolution of law in the European context shows how having access to law is a relatively new ideal.⁹⁸ From a citizen's point of view, better access to the law reduces costs – in the form of legal fees – and enables them to be informed. From a legislator's point of view, making sure people know the law (or at least have access to it) can help compliance with it and provides people with access to the government services that they may need. On the other hand, when citizens are excluded from their rights because of difficulties to access and understand law, the mandate of the rule of law can be called into question.

In Section 4, all examples listed bar two originated from state institutions; only *DoNotPay*, UK, was from a private individual, and *ElectKB* from two Australian universities. While it is true that government bodies and institutions like universities might be more prone to develop projects for better accessibility to the law, private projects do also exist. A popular example thereof is

⁹⁶ The Economist, 'A benefits scandal sinks the Dutch government' *The Economist* (Jan 23).

⁹⁷ L. Henriques-Gomes, 'Robodebt: total value of unlawful debts issued under Centrelink scheme to exceed \$1bn' *The Guardian* (Jun 10).

⁹⁸ T. Herzog, *A Short History of European Law: The last two and a half millennia* (Harvard University Press 2018).

ROSS. ROSS (built using IBM Watson) accepts legal questions in plain English and returns answers based on legislation, case law, and other sources.⁹⁹

To conclude, notwithstanding the sponsor of an automatically processable regulation, projects aiming at better accessibility have enabled citizens to better query the law, use information and argument retrieval systems, check their eligibility for a wide range of services, and compute benefits that they are able to obtain. We see thus the promise of more accessible law more as one more closely achieved while raising fewer problems than the one of higher efficiency in legal processes.

5.3 Is automatically processable regulation compatible with bureaucracies?

Our typology-based review of projects in the domain of automatically processable regulation shows that there are few projects being followed by states which have a high degree of mediation by computers. This stands in contrast with the state of research highlighted in Section 2.2 and begs the question of how to explain this gap. One explanation can be that, at this stage, convincing policymakers of the value of automatically processable regulation projects is difficult.¹⁰⁰ The fact that the public service is able to deliver on its mandate with its current (mostly manual) setup partially explains the lack of pressure to change what works, even if efficiency gains could be realised. Adding to this mix is that it is politically a difficult selling point to seek to disintermediate humans, as any politician standing for re-election would fear alienating both those who could lose their jobs, and those favouring human contacts over using technology.

In addition, many public services have not sought to model the costs and benefits that could be derived from turning regulation into automatically processable regulation. In our research, we could not come across anyone who even attempted such a cost-benefit analysis. It is hence difficult for any decision-makers to envisage what trade-off this could entail. Moreover, many projects yield implementations with accuracies lower than 100%, even if some academic research have achieved high scores. Anything short of 100% will bring further difficult questions of explainability and accountability for what occurred inaccurately. Again, as long as it is not possible to measure (economically, but not only) the extent of consequences when inaccuracies of implementing automatically processable regulation come about, it will not be possible to weigh it against the benefits. Decision-makers, including citizens, need to be able to gauge the extent of worst-case scenarios. Without the possibility to do so, well-placed and misplaced fears of algorithms introducing a high degree of disintermediation of humans will remain.

5.4 In which context should automatically processable regulation be applied?

There is, however, yet another and simple explanation for the lack of projects with high degree of disintermediation of humans: Certain statutes are more prone to being encoded by humans than others. In fact, the question of what constitutes a difficult legal issue to turn into an algorithm was already taken up in 1987 by Gardner,¹⁰¹ who devises a test by asking a series of yes/no questions (such as, whether common-sense-knowledge rules could provide an answer to a legal question, whether the issue is a positive instance of what the predicates describe, or a negative instance). In this context, our surveyed projects show that not all legislation (or parts of legislation) are suitable to be processed automatically. *LabPlus* in New Zealand¹⁰², during their three-week sprint, identified five characteristics of legislation (or parts thereof) for which they see value in transforming it into legal code and we can relate these to the other aforementioned projects: (1)

⁹⁹ Ashley, *Artificial Intelligence and Legal Analytics: new tools for law practice in the digital age*.

¹⁰⁰ Interview with Scott McNaughton, March 8, 2021.

¹⁰¹ A. v. d. L. Gardner, *Artificial Intelligence and Legal Reasoning* (The MIT Press 1987).

¹⁰² LabPlus, *Better Rules for Government: Discovery Report*, 2018.

legislation that involves calculation (such as *Mes Aides* in France); (2) legislation that involves a process that requires factual information to determine application, eligibility, entitlements, or coverage (such as *ElectKB* in Australia); (3) legislation that prescribes a process that is used repeatedly (such as for vessel registration in Canada); (4) legislation that prescribes a compliance process or obligation (such as traffic regulations embedded in autonomous vehicles); and (5) legislation that prescribes a process or a system that can be delivered digitally (such as *DoNotPay* in the UK).

This motivates the question as to whether attempts will move to other areas of law – such as criminal law or open-textured legislation. But at the moment, in terms of implementation, we are still far from it, and rather at the stage of policymakers convincing key stakeholders of the added value of encoding law. The New Zealand research also acknowledged – in a reference to certain limitations such as open-textured terms or ambiguity – that ‘it is difficult to produce machine consumable rules if the policy and legislation has not been developed with this output in mind’.¹⁰³ Yet, they show that their approach of bringing multidisciplinary teams of ‘policy analysts, legislative drafters, service designers and software developers’ in order to commonly design legal code ‘taking a user-centric approach that focuses on how the service could most effectively be delivered’ can lead to workable solutions that can effectively be implemented into practice.¹⁰⁴

Furthermore, with humans mediating the encoding of regulations, looking at technical limitations only would be too narrow. The tax code is arguably one of the most computation-intensive domains of law. It is therefore informative to look at roadblocks that researchers have encountered when trying to encode it. In 2005, the Danish Tax authorities introduced a digital tax collection system that needed to be dismantled ten years later after errors emerged which meant that the state had foregone collecting close to USD 14 billion in taxes.¹⁰⁵ One of the key points around governance was that ‘data and systems often were so badly documented that [the tax authority] did not have sufficient insights into them’.¹⁰⁶ This illustrates that even in a context where automatically processable regulation might emerge, its success is far from certain. Denmark’s parliament understood the lessons well (not only from the tax scandal but from a fraud scandal building on automated processes¹⁰⁷) and thus adopted seven principles for ‘digital ready legislation’.¹⁰⁸ These lessons-learned show that even if all the possible technical limitations were removed, many other complex parts would need to align for the system to work.

5.5 Deliberation on Typology Dimensions

When designing the proposed typology, we pondered on different dimensions to include or not. Three specific dimensions were heavily discussed and, in the end, left out: *efficacy* (instead of accessibility), *stakeholder groups* (instead of potential for divergence of interests), and *determinism* in automatically processable regulation. In the following we justify our choices.

First, regarding *efficacy*: Via automatically processable regulation, it is possible to argue that the law becomes more effective. Those who are eligible for benefits become aware that they can seize upon these rights; those who wanted to engage in anti-social behaviour become aware of the disincentives for doing so. Yet, both applying to benefits as much as not engaging in an otherwise illegal act are not the primary-order consequences of making regulation automatically

¹⁰³ Ibid.

¹⁰⁴ Ibid.

¹⁰⁵ M. Barrett, ‘After years of scandals, here’s how Denmark’s reformed tax authority will look’ *The Local DK* (Jun 14).

¹⁰⁶ B. Alfter, *Automatic Society 2019: Denmark*, 2019.

¹⁰⁷ D. Segal, ‘Where in the World is Denmark’s \$2 Billion?’ *The New York Times* (Oct 5).

¹⁰⁸ Danish Agency for Digitisation, ‘Seven principles for digital-ready legislation’ (2021) <<https://en.digst.dk/policy-and-strategy/digital-ready-legislation/guidances-and-tools/seven-principles-for-digital-ready-legislation/>> accessed Jun 23.

processable, but only the secondary order. The primary order is that individuals become *aware* of the law, and as such are then able to take up their rights or remain in the legality. This first order of *accessibility* to the law is much more tangent and direct; the one of *efficacy* of the law much more uncertain and dependent on a range of questionable assumptions (e.g. whether those engaging in criminal behaviour think rationally, or the difficulty of testing deterrence when no action took place).

Second, regarding *stakeholders*: There is value in mapping out the different stakeholders of each individual project on automatically processable regulation, notably in mapping out specific stakeholders in order not to fall prey to the myth of the monolith – even within institutions, units wage turf wars on power and resources. However, such an exercise can make it difficult to compare projects without otherwise lumping entities together. The problem of lumping entities together is that it makes void the reason of attempting to map out stakeholders in the first place, namely, to tease out where their interest could clash, and to deduct therefore whether this could either impede implementation all together, or lead to many further issues. This is why we assess the value in gauging the degree of *potential for divergence* in being much greater than in merely differentiating between different sponsors, implementers, beneficiaries, and users.

And third, regarding determinism: The (legal) literature on the topic of automatically processable regulation often suggests a distinction between ‘deterministic’ (sometimes referred to as ‘rule-based’) and ‘non-deterministic’ (e.g., ‘probabilistic’) algorithms.¹⁰⁹ While this distinction seems a compelling proxy for our proposed dimension of ‘mediation by computers’ at first sight, it is – depending on its concrete usage – technically imprecise or even incorrect: First, the term ‘algorithm’ refers to the *unambiguous* processing of input and return of an output within a finite amount of time¹¹⁰ and, thus, typically already refers to *deterministic* algorithms, where the output of the algorithm is completely determined from its inputs. Such algorithms behave like specified mathematical mappings of inputs to outputs, and applying the same input always leads to the same output.¹¹¹ This is not true for *non-deterministic* algorithms, where applying the same input twice might lead to different outputs. Importantly for the discourse on techno-regulation, however, an algorithm being *non-deterministic* does not imply that the output of that algorithm is ‘random’ or ‘arbitrary’. Rather, it is merely the *execution* of the algorithm that behaves in a non-deterministic way while the output is valid according to the specification of the algorithm (an algorithm that returns outputs which violate its specification is called *incorrect*). Important categories of non-deterministic algorithms are *probabilistic* algorithms, where the non-determinism stems from the introduction of randomness into the execution, and *concurrent* algorithms, where the non-determinism stems from the sequence or timing of events in the execution. However, in both cases, this non-determinism in the execution does not imply that the algorithm will produce arbitrary outcomes, and we therefore propose that this distinction is not useful when referring to the high-level implications of automatically processable regulation.

6 Conclusion

Understanding what real-world applications of automatically processable regulation exist and how they work is a difficult undertaking. Predicting the impact of such attempts is even harder.

In this article we review multiple projects that fall under the heading of automatically processable regulation, what others have also called computable law. The analysis reveals that multiple dimensions exist and interact with one another. In order to enable comparability of

¹⁰⁹ See notably for the distinction Ashley, *Artificial Intelligence and Legal Analytics: new tools for law practice in the digital age*.

¹¹⁰ Math Vault, 'The Definitive Glossary of Higher Mathematical Jargon'.

¹¹¹ Hans Delfs and Helmut Knebl, 'Probabilistic Algorithms', *Introduction to Cryptography. Information Security and Cryptography* (Springer-Verlag 2002).

projects, we propose a typology to classify ongoing and future implementations of automatically processable regulation. In our typology, we draw on three distinct dimensions: first, the primary scope of an automatically processable regulation project, second, the potential for divergence of interest within a project, and third, the degree of disintegrational of humans in a project.

The first dimension asks primarily who the main beneficiary of a real-world implementation of automatically processable regulation is. If the focus rests primarily on laypeople, such a project is considered to be focused on accessibility to the law, otherwise, the primary focus rests on efficiency. The different aims are key with respect to the promises attached to different projects of automatically processable regulation as well as the assessed reachability of such promises (see Section 5). In fact, the real-world applications analysed throughout the article challenge the promise of higher efficiency. The goals of projects focusing on better accessibility to regulation were achieved with fewer problems by providing laypeople with tools to demand benefits or challenge ordinances.

The second dimension captures the multitude of stakeholders when it comes to implementations of automatically processable regulation and the differences of interests among them. This dimension indicates the complexity of automatically processable regulation when it becomes more than just a thought experiment of technical and legal scholars (especially when implemented in the midst of tangled bureaucratic process). Thus, even in instances where automatically processable regulation might seem technically feasible, the human element and interaction with stakeholders should not be underestimated, but thoroughly considered before pursuing a concrete implementation.

The third dimension looks at in how far humans are taken out of the loop when a specific project of automatically processable regulation is implemented. This disintermediation occurs during implementation and applies to both procedural elements in the code itself and the extent to which the data is being curated. The degree of mediation by computers is highly dependent on different factors, such as the complexity of the laws being encoded, the actual code used for such processes, and the reliance on (past) data (depending in turn on the degree of curation of any datasets used).

We argue that the systematic classification of automatically processable regulation projects is important for at least three reasons. First, it allows to *compare projects*. The comparability is important to develop a holistic understanding of what automatically processable regulation entails, meaning understanding its impact on the legislative process, individuals, and society at large. Second, the systematic approach provided by the proposed typology enables a *nuanced discussion on the issues* resulting from automatically processable regulation. Importantly, the discussion of the surveyed projects shows that the primary aims of automatically processable regulation vary vastly and the potential of achieving its promises does as well. Unbundling and disentangling these issues is important to not dismiss projects that have inherent merit by falsely tying them to ones that are inherently unachievable. Finally, the typology *enriches the policymaking discourse*. Comparability enables drawing lessons from different projects, not only within the same jurisdiction but as shown within this article from across the planet. It is then possible to draw policy lessons from these projects, which is especially key in this still nascent field that is gaining attraction by academics, the media, and policymakers. For instance, it can be remarked how projects with higher disintermediation are geared towards efficiency, that no project implements the highest forms of disintermediation, and that there is therefore a gap between academia, the media, and policymakers. Guiding the policymaking discourse at an early stage with such a nuanced approach is key to ensure success of projects and to, in warranted cases, discount projects based on prior lessons within other jurisdictions.