Scoping AI & Law Projects: Wanting It All is Counterproductive

CRCL23: Computational "law" on edge

DENIS MERIGOUX, Inria, France

1 INTRODUCTION

In their recent presentation, Gebru and Torres [2023] relate how Gebru, originally an electrical engineer, came from criticizing technical aspects of machine learning models to discovering and making explicit the underlying ideologies that drive the latest advances of machine learning. The leaders of the field of Artificial Intelligence (AI), they say, are influenced by the TESCREAL[1] bundle of ideologies. These ideologies view machine learning as a key enabler to artificial general intelligence (AGI), a god-like entity that can bring both utopia and apocalypse to mankind. While the debate around AI and its future is raging and goes beyond the scope of this paper, the presentation of Gebru struck me particularly as I share the externality of her point of view. Indeed, I was trained in formal methods–a subfield of computer science close to its theoretical foundations and aimed at raising the level of assurance of critical software–and entered incidentally the intersection of Computer Science (CS) and Law. This intersection is dominated by the field of AI & Law, as per the names of the eponym flagship journal and main conference (*International Conference on Artificial Intelligence and Law*, ICAIL).

Why is AI dominating this intersection? As new academic venues like ACM's CS & Law conference, the Programming Languages and the Law (ProLaLa) workshop and the CRCL conference and journal are emerging for scholars at the intersection of CS and Law, the consensus around the affiliation to the broader field of AI is being questioned. In this paper, I will investigate what the affiliation to the field of AI means, ideologically and technically, for legal technology research. First, I will make a link between the self-proclaimed goals of leading pieces of research in AI & Law and two powerful ideologies that aim at defining the ideal form of government, legal formalism and cybernetics. This link will help me emphasize the technological goals pursued by the field of AI & Law, and why they create a tension with the rule of law and democratic decision-making as explained by Hildebrandt [2020]. Second, I will analyse how the overreaching goals of AI & Law have constrained the technological solutions put forward by the researchers into systems that "want it all". Finally, I will argue that pursuing systems that "want it all" is counterproductive both technically but also socially as it delays the adoption of advanced legal technologies in domains where it is really needed.

This position paper aims at sparking a debate in the community and maybe shuffle the priorities of the field. More research will be needed to fully validate the claims that I make here, and I am sure valid opposing arguments could be raised. But my personal position expressed in this paper is that rather than frantically looking for a way of integrating the latest glowing artifact of AI like ChatGPT into the existing AI & Law agenda, I would suggest that the field comes back to the roots of the scientific method and starts building better evaluation frameworks for the output of its research. By adapting its technical solutions to the imperative of the rule of law and the real needs

[1]Transhumanism, Extropianism, Singularitarianism, Cosmism, Rationalism, Effective Altruism, Longtermism.

of the legal profession, a new era of applied research at the intersection of CS and Law could be ushered.

2 THE TANTALIZING PROMISE OF GOVERNMENT BY MACHINES

In this section, I will take the leading scholars of AI & Law at their own words and try to make some links with broader ideologies that share an ideal vision of government by machines. Conveniently, the AI & Law journal just has published a retrospective of its three decades of existence in the form of the articles by Governatori et al. [2022], Sartor et al. [2022] and Villata et al. [2022]. However, this retrospective does not bring a reflexive point of view on the field and merely states which technical solutions have been brought about by the community, not how the community had selected the problems to solve.

Hence, I must look deeper into the position papers and opinion pieces of the leading scholars of AI & Law to understand *why* they chose to tackle the formal representation of legal knowledge of court decisions prediction. A first stop is the white paper by Genesereth [2015] that describes the orientation for research done at CodeX, the Stanford center for legal informatics. At first, Genesereth starts with a real-world problem: laws can be very complex and that affects compliance, efficiency and trust. But then, without considering how actual lawyers and subjects of law deal with legal complexity, he goes for textbook techno-solutionism: "*fortunately, these problems are not insurmountable. To the extent that they are information problems, they can be mitigated by information technology*". The technical solution put forward here is computational law, which basically amounts to build a formal representation of legal knowledge and have legal decisions computed from it. He then lists examples where this has been put in practice, and acknowledges a fundamental limitation: "*the resolution of [the problem of the open texture of laws] is to limit the application of Computational Law to those cases where such issues can be externalized or marginalized*". But this limitation does not seem to prevent computational law to be deployed ubiquitously: "*you are walking through the woods of Maine and see an attractive flower. You take a photo with your iPhone. Your plant app identifies it as a type of orchid and lets you know. At the same time, your legal app tells that, no, you may not pick it*". Finally, going further than the ubiquitous use, Genesereth claims that "*in a way, Computational Law is the next step in the evolution of the legal system*". The discourse here ends on a messianic note and computational law steps up from a technical solution to real-world problems to the inevitable future of the legal solution.

When looking closely, we can see that computational law as touted by Genesereth is heir to the old ideal of legal formalism. In our context, the important aspect of legal formalism is that laws and their enforcement should follow logical rules and must not depend on contextual elements that can change over time and place. Legal formalism is an old idea that peaked at the junction of the XIXth and XXth centuries and was already back then criticized, for instance by Dewey [1924]. But the field of AI, which was dominated at the end of the XXth century by the formal representation of knowledge, quickly recognized that its goals coincided with legal formalism: "*AI and Law is much more than an applications area. Its concerns touch upon issues at the very heart of AI: reasoning, representation, and learning. For the AI researcher interested in symbolic methods—or methods of whatever stripe—that are focused on providing explanations and justifications, AI and Law is an excellent arena*" [Rissland et al. 2003]. Thus, the general promise of efficiency, accessibility and uniformity of legal formalism are the primary motivations of legal knowledge representation projects. Particularly, the effort to standardize the encoding of legal texts and rules in LegalRuleML weaves these promises in the context of the Web: "*legislators, legal practicioners, business managers are, therefore, impeded from comparing, contrasting, integrating, and reusing the contents of the texts, since any such activities are manual. [...] In doing so, the general norms and specific procedural rules* *in legislative documents, the conditions of services and business rules in contracts, and the information about arguments and interpretation of norms in the judgements for case-law would be amenable to such applications*". [Athan et al. 2013]. The formal representation of legal knowledge and concepts is thus viewed as the gateway to finally access the benefits promised by legal formalism, though this intervention of AI into Law and society in general gathered criticism early for its tendency to conflate judgment with computation, as per the thesis of Weizenbaum [1976].

The specific uses for formal representations of legal knowledge are diverse: most AI & Law scholars state that the intended use is to explain, justify, teach or better understand legal concepts. However, once a formal representation is made, it can be used for more than that. In that sense, the applications of AI & Law technologies can be qualified as opportunistic; once created under benign motivations, an artefact can be repurposed for more involved applications. I will give here three examples of this potential switch in applications for a given technology. First, the question of whether judges should be replaced by the computer execution of a formal legal knowledge representation has been asked in the past by D'Amato [1977]. Some AI & Law scholars have rejected this idea categorically: "*A juridical machine can thus only be an aid to the jurist and not a substitute for him. We shall have no 'electronic judges' in the world to come, any more than we shall have a machine to rule us*" [Mehl 1959]. But recently, replacing judges with machines has taken a very concrete outlook with the advent machine-learning based prediction [Medvedeva et al. 2023]. As a good illustration, the prediction tool for COVID-affected exceptional business rent exemptions by Parton et al. [2024], officially supported by the Italian government, is first presented as a tool for explaining court decisions, but then the authors note that it could be used as an assistant for judges, even though nothing prevents the judge from simply taking the tool's output as their final decision. Second, a classic use for formal representations of legal knowledge is assistance to legal drafting: "*an executable, logic-based representation of rules and regulations can be used not only to apply the rules, but to aid the process of drafting and redrafting the rules in the first place–a point that was made by Allen [1956]*" [Sergot et al. 1986]. But today the "assistance" to legal drafting is switching to a madatory requirement that laws should be drafted in a "digital-ready" fashion, with explicit guidelines about how the legal text should follow a precise logical structure. Indeed, Danemark passed its Digital-Ready Legislation Act [Plesner and Justesen 2022], and the OECD [2019] and the European Commission[2] are supporting similar initiatives. Third, the rise (and fall) of blockchain technology and its derivatives made the idea of automating contracts completely more realistic, as noted by Crafa et al. [2023]: "*Since parties are free to express their agreement in the language and medium they choose (freedom of form, a principle shared by modern legal systems), drafting a contract by using a programming language (rather than, as usual, natural language) seems a valuable option. Advantages are in terms of speed-up, lack of ambiguity, and automatic and transparent enforcement of the contractual clauses. For this reason, several projects are being developed for defining programming languages to write legal contracts, e.g. [Basu et al. 2019; Contributors 2018; Dwivedi et al. 2021; Foundation 2019; He et al. 2018; Merigoux et al. 2021; Wright et al. 2019]*"[3].

The tantalizing opportunity of using AI & Law applications beyond academic circles to directly intervene in public policies is where AI & Law meets a second powerful ideology: cybernetics. The science of systems founded in the second half of the XX$^{\text{th}}$ century was quickly endorsed

[2]See the "Digital-Ready Drafting" track at the SEMIC 2023 conference.

[3]Funnily, I have personally been approached by three different groups of people asking me how I could formalize into an executable program the governance contract defining the rules of their Decentralized Autonomous Organization (DAO). I had to go to great lengths to explain to them that it was a bad idea.

by scholars that wanted to apply it to society and the State[4]. But before that, in France, Mehl [1957] theorized how the State apparatus could benefit from cybernetic principles to increase the efficiency of its administration: "*the administration can thus be seen as a cybernetic system, but with its own specific aspects. Administration operates solely on information. Its counterpart in the world of machines is the 'computer', not the machine tool. Administrative information is rarely imprecise and sometimes erroneous. Administrative action is altered by random phenoments. As a result, it takes on the appearance of a a strategy*"[5]. Later, Mehl and Breton [1970] and Catala et al. [1974] created of one the first French legal databases in order to materialize the cybernetic ambitions of automating the analysis of legal cases by the administration [Mehl 1959]. Knapp and Vrecion [1970] corroborates this link between cybernetics and the nascent AI & Law techniques, and provides more examples of similar projects in the USSR, Czechoslovakia and the United States. In Italy, Contissa et al. [2021] relates the theoretical work of Frosini [1968, 1973] inspired by the same ideas. More recently, cybernetics is has been making a comeback in AI & Law, for example through the work of Bourcier [2017], Mehl's former student, Potvin [2023]; Potvin et al. [2021] and their idea of a computer infrastructure for distributing rules to actors, or Sileno [2016], that cites a couple cybernetician references and explicitly takes a cybernetics approach: "*[...] in order to fulfill their mandate, the responsible authorities must put in place adequate activities to target known and hypothetical non-compliance patterns, along with anticipatory, discovery mechanisms to unveil new ones. But non-compliance is only half of the story. [...] In short, public administrations have to adapt their allocation of resources and scheduling of activities in accordance with the social environment in which they operates, and to the requirements set by the legal system*". The crux of the issue to integrate AI & Law with cybernetics is to make a precise enough formal model of law and society as a system to be able to analyze it. This modelling activity involves formalization and knowledge representation, which is why the technological solutions developed in the context of AI & Law are highly relevant for cybernetician endeavours.

So far, we have emphasized the links of the AI & Law community with two powerful ideologies: legal formalism and cybernetics. These links all point to a shared goal, which is building models or formal representations of law and the objects it regulates, in order to first explain and justify existing phenomena, and maybe later to automate aspects of the legal system and the administration of society. While AI & Law projects are careful when stating their ambitions and most scholars would reject the literal idea of government by machines when asked about it, it becomes evident that the field is attracted to an utopia where law is formally codified and instantly accessible through digital mediation, automatically enforced and applied without bias and enables a very efficient society where uncertainty, risk, frictions and delays are reduced to a minimum. I will not discuss here whether that utopia is desirable or not; however I can remind some of the contradictions between this utopia and the rule of law as it exists currently in democratic societies. I would refer the reader to the formidable work of the COHUBICOL typology [Diver et al. 2022] for more analysis and examples about this. Personally, I believe these contradictions are best evident when analysing the proposition of Fraser [2021] and de Sousa and Andrews [2019] based on their experience advising the Australian and New Zealand government: for better efficiency and automatic enforceability, a "rules as code" [Mohun and Roberts 2020] or "better rules" approach should be adopted. The core of this approach is "*[...] the use of the multidisciplinary team that includes people skilled in policy, legal, business rules, drafting, programming and service design, working together in an iterative fashion to*

[4]The infamous Cybersyn experiment in Allende's Chile [Medina 2006] showed how information technology could power the new infrastructure required to control a whole national economy based on rational algorithmic principles and control theory with feedback loops.

[5]My translation from French.

develop rules" [Fraser 2021]. Here, "rules" is used in a very broad sense, as the presentation purposely blurs the institutional divisions that comes with the rule of law. "*By collocating with drafters and coders, this group can then simultaneously co-draft human and machine readable versions of the rules for testing — with humans and machines. This allows for more holistic modelling of impacts, and provides and the opportunity to test the coded rules with end users (regulated entities, service providers, etc.) before publication*" [de Sousa and Andrews 2019]. Inspired by the Agile [Beck et al. 2001] methodology for producing software, this co-drafting and co-implementing of the rules (legislation) basically wants to put together the lawmaker, the administration and the subjects of legislation in a room to develop together the law and the way it will be enforced. This proposal actually goes against centuries of political philosophy where separation of powers is considered as a core pillar of the rule of law. Indeed, what happens when the interests of the lawmakers, administration and subjects of the legislation are not aligned? The separation of powers and deliberative assemblies allow diverging interests to be reconciled with compromises while avoiding conflicts of interests, something the "rules as code" or "better rules" approach ignores completely (or claims to solve with complete transparency of the source code implementing the rules). I do not deny that the separation of powers and current institutional schemes can entail some inefficiencies, but I personally believe that these do not warrant throwing the rule of law out with the pivot to more digital-friendly practice.

Hence, the technologies developed in the field of AI & Law share the common goal of modelling law and society to get closer to the benefits of the tantalizing promise of government by machines. What are the (practical) results since the 1980's? So far, not much; as pointed already by Oskamp and Lauritsen [2002] who justify the lack of success by two reasons (summarized). First, law is too complex and hard to model. Second, the users (lawyers) are too conservative and illiterate about technology. "*For most of the past fifteen years practicing lawyers and AI researchers appear to have been locked into parallel worlds of theoretically uninspiring implementations and tiny brittle research applications. Robust traffic across that disciplinary divide has yet to develop.*" That was true in 2002 and, I claim, to some extent still true in 2023, so perhaps there is a deeper explanation behind the two reasons provided here.

3 TECHNICAL SYSTEMS THAT WANT IT ALL

I claim here that because of the utopian nature of the goals set by the field of AI & Law to itself, it has been compelled to only research technologies that are all-encompassing, attempting to solve every problem at once while twisting the practice of the rule of law. In a nutshell, AI & Law technical systems "want it all" and it is the cause of their lack of success in the real world. To back up this claim, I will discuss technical aspects based on the AI & Law recent retrospective by Governatori et al. [2022], Sartor et al. [2022] and Villata et al. [2022].

In the first decade, the focus of AI & Law was discovering the logical features necessary to formally represent laws, norms, regulations, court cases and argumentations. This focus yielded several key theoretical results (deontic logic, defeasible logic, argumentation schemes, isomorphism) that constitute an efficient state of the art for legal formalization. However, and since the beginning, the confrontation between theory and practice yields poor results: "*Sergot et al. [1991] argued that it is appropriate to follow an isomorphic approach if the legislation is itself well structured, but otherwise this approach might become cumbersome. However, legislation is very often not well structured. In such a case, isomorphism would lead to a poorly structured knowledge base, one which fails to correspond to the 'real world' problem.*" [Governatori et al. 2022]. The crux of the issue is that law does not always conform to the expectations of legal formalism: some enacted and enforced norms may be contradictory or ill-conceived, but they are nonetheless in force. This problems is solved through social processes, where parties chose an interpretation that suits their interest. A vague or ambiguous regulation can even be made on-purpose to serve the interest of the regulator [Torny 2005]. Hence, AI & Law projects tackling real-world legal situation with knowledge representation tools are forced into a dilemma, that I will illustrate with quotes from the seminal AI & Law paper by Sergot et al. [1986].

(1) Either developers simplify the model to avoid the complexity of the real-world and stay within the realm of what is formalisable but then the model loses a lot of its utility: "*the simplest way to handle vagueness is to assume that the vague concepts always apply and to use this assumption to generate qualified answers*".

(2) Or developers complexify the model by incorporating more and more bits of real practice into the model. However this may require escalading weird logical features to account for the irrationality of reality, at the risk of rendering the model unusable (impossible to execute or maintain it): "*a more sophisticated approach might combine this with the use of rules of thumb that reduce vague concepts to concrete ones, but are not guaranteed to cover all cases. The rules of thumb arise from the analysis of previous cases. We deliberately avoided such complications and chose the simpler alternative in our implementation of the act*".

Because of the influence of the ideologies discussed in the previous section, AI & Law scholars need their model to be as precise and rich as possible to preserve their potential for infinite reuse into all areas of applications. Hence, they usually want to choose option (2) of the dilemma, unless they hit the hard technological limitations of their tools, in which case they stick with option (1). Consequently, most AI & Law scholars tend to choose technological platforms that allow for very general and open-ended computer modelling. During the first decade of AI & Law, at the turn of the 1990's, the most versatile formal modelling tool around was Prolog [Colmerauer and Roussel 1996], which is not surprising as the goal of Prolog was to be the ultimate meta-language in which to declare formal systems. But modelling versatility in computer science tools comes with trade-offs. For example, running Prolog programs requires the use of a Horn clauses solver that requires a heavy runtime and may be a source of inefficiency. While it is possible to write efficient programs in, for example, SWI-Prolog [Wielemaker et al. 2012], the efficiency is conditioned by the use of a very strict subset of the features of Prolog that is correctly optimized by the interpreter, a subset that may not match what is required to elegantly and concisely model the law. Moreover, the diversity of opinions and research projets around Prolog transformed it into family of languages [Körner et al. 2022, Tables 1 and 2] that share a common core, but where each has its strengths, weaknesses and quirks. Furthermore, there isn't yet a consensus in the AI & Law community about the exact variant of Prolog to be used. For instance, new AI & Law projects by Arias et al. [2021] or Lim et al. [2022] have switched to using Answer Set Programming (ASP) or its variants, ASP being itself a variant of Prolog with a completely different semantics and set of implementations.

The diversification and non-interoperability of models built by AI & Law projects attracted the attention of another branch of the AI research community, focused on semantic representations and ontologies. This shift corresponds to the second decade of the AI & Law retrospective, and at first the goal of introducing ontologies is well circumscribed by Breuker et al. [2004]: "*an ontology makes explicit the concepts and their properties one is committed to in modeling a domain. Note that we do not consider an ontology itself to be a model of a domain: it is used to have unambiguous and shared terms in the model*". So then, ontologies could be viewed as a soft tool for aligning diverses models into interoperability. As such, they were object of great developments in the 2000's and 2010's. For instance, Barabucci et al. [2010] and Palmirani and Vitali [2011] designed Akoma Ntoso, an ontology for structuring the presentation of legal documents around the world, which then became an UN-endorsed [Peroni et al. 2017] standards adopted by many legal publication offices around the world. The success of this ontology relies on its very low semantic content: its killer features are the ability to declare precisely what is a paragraph, a list, a document, etc. and cross-reference these items. Akoma Ntoso as an ontology does not try to express what the law means, but merely how it is structured. Thus, it falls short of actually modelling the law in the sense of the ideologies discusses in the previous section.

But more generally, as a true dual-use technology and a very versatile modelling tool, ontologies can also turn themselves into the formal models they were meant to align: "*as ontologies contain generic knowledge, cost-effective knowledge engineering may benefit from its reuse potential. Indeed, one can argue that the use of ontologies in AI comes from research in the late 80s and 90s that aimed at improving knowledge engineering by creating 'well-structured' knowledge bases that would not only solve the problem at hand but be more maintainable, easier to extend, etc. In this sense, ontologies are then very much an engineering tool. This role of ontologies implies the use of an inference engine that is used to conclude specific goals*" [Breuker et al. 2004]. Here, the tantalizing opportunity to expand the use of ontologies in the legal domain into more computationally involved uses thrived: "*there is urgent need to find a robust and expressive XML annotation, compliant with the Semantic Web technologies, able to meet all the unique particular aspects rising from the legal domain and in the same time close the gap between legal text descriptions, using XML techniques, and norms modeling, in order to realize an integrated and self-contained representation of legal resources available on the Web*" [Palmirani et al. 2011][6]. One of the most salient work in this line of research led to the creation of the LegalRuleML ontology [Athan et al. 2015] on top of Akoma Ntoso, which embarked all the logical features needed for legal formalisation discovered in the previous decade. In a *grand finale* proof of concept, Palmirani and Governatori [2018] present the integration of all the ontology-based tools to showcase an example of a legal model of GDPR being used for automatically enforcing it (on a specific platform). However, this proof of concept is criticized by Novotná and Libal [2022]: "*[...] they do not deal with multiple interpretations and they do not specify the cooperation with legal experts. Secondly, they don't provide use examples and any evaluation of correctness or usability of the system.*"

As Figure 1 depicts, the actual technological artefact behind this *tour de force* is made of a dozen different software tools interacting with each other, making up a complex architecture that required years of infrastructure building, typical of the high-modernist (in the sense of Scott [1998]) approach to computer science that I criticize in my PhD dissertation [Merigoux 2021]. "*In our framework, presented below, we find LIME and RAWE [Palmirani et al. 2013], which are two web editors (JavaScript) capable of semi-automatically marking up the text in Akoma Ntoso and the manually formalized norms in LegalRuleML. PrOnto [Palmirani et al. 2018] is a legal ontology for modelling GDPR concepts and axioms. It feeds concepts and predicates to the legal rule-modelling layer in order to make the formalization consistent and harmonized. Regorous [Governatori 2015] is a tool (written in Java) that makes it possible to design BPMN 2.0 and to connect each step of the process with the legal rules. Regorous provides an API to SPINdle [Lam and Governatori 2009], a defeasible legal reasoning engine. Regorous presents at the end the results of compliance checking in a user interface for the end user*" [Palmirani and Governatori 2018]. The problem with this approach is that in the process of creating this integrated platform for modelling GDPR and its evolutions in time to enforce it automatically, the authors have ended up reimplementing a whole dedicated software engineering toolchain that relies on *ad-hoc* tools that do not interoperate with standard software engineering tools. The ontology inference engine used to actually execute the LegalRuleML annotations may entail a

[6]A similar quote can be found in [Palmirani et al. 2009].

significant refactor of an existing IT system in order to be used in real-world applications. The LegalRuleML rules are themselves a *de facto* form of source code that can only be edited and viewed through custom editors and visualizers, that may not support all the features of modern standard code editors. For instance, there is no version control system for this source code (though the rules themselves track code of the versions of the law)[7]. Ironically, the ontology architecture made for aligning legal standards goes against the Unix philosophy that underlies a lot of modern software engineering best practices: "*Write programs that do one thing and do it well. Write programs to work together. Write programs to handle text streams, because that is a universal interface*" [Raymond 2003]. Simply put, integrating the toolchain of Palmirani and Governatori [2018] in an existing IT system is highly disruptive technically and its ability to scale up is not yet demonstrated.

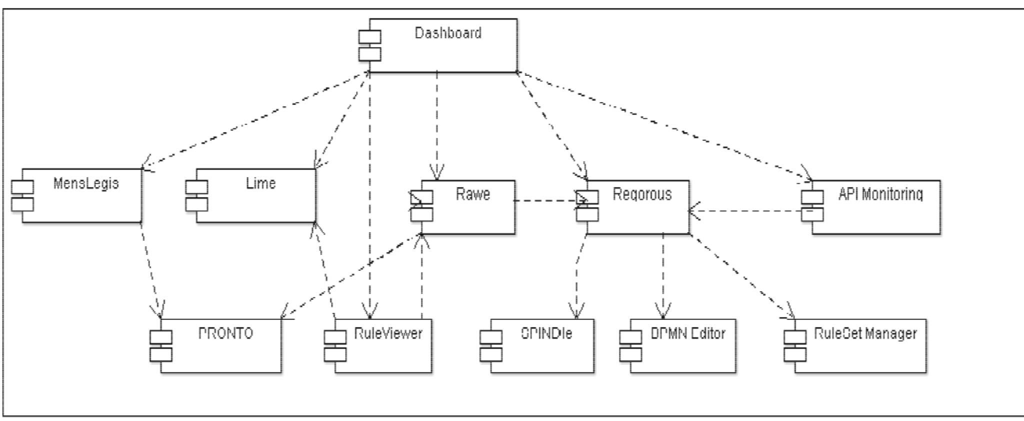

Fig. 1. Legal Tools Architecture, quoted from [Palmirani and Governatori 2018]

But this paradox is not surprising, as all big ontology-making endeavours are doomed to becoming more and more integrated and monolith-like as they approach actual real-world usage[8]. This was already mentioned by Breuker et al. [2004]: "*An interesting problem that arises is the introduction of an inference bias. Valente et al. [1999] show that ontological choices are strongly influenced by the purpose of the ontology. That is, the same knowledge will be structured or formalized differently depending of how it will be used by the reasoner in reaching the desired conclusions in a specific context. This indicates that reusability is a good idea, but it can never be accomplished completely*". This critic is put in even clearer terms by Brewster and O'Hara [2004]. I add here two points to the argument. First, that the inference bias also translates to a technological bias for ontologies as software artifact that tend to build separate ecosystems from the rest of software engineering, preventing them from harvesting the scale-up benefits coming from the huge infrastructure investments made in standardized software engineering tooling. Second, that the technological bias leads to

[7] I have personally been able to inspect the core rules infrastructure of a major payroll editor. As they went with a no-code approach based on rules modelling not unlike Palmirani and Governatori [2018], they put all their rules in a database without versioning them. Now they have huge business problems with maintaining their system as they cannot use branches, merge requests, continuous integration but instead only have a "dev" and "prod" version of the rules database.

[8] Personally, I experienced this during an internship in 2015 at a defense contractor that tasked me with building an ontology to interoperate different message text formatting systems in use by army equipments. As I was writing a message translator using SPARQL queries, I suffered from the poor software engineering environment and the inability to interoperate with anything. After I left, nobody used my working proof of concept and stuck with using an XSLT-based transformation engine with hand-written XSLT translation rules. The hand-writing of the rules continued to burn-out dozens of engineers per year long after I left.

a bias in the users of the technological tools: rather than being usable directly by either lawyers or programmers, the tools require both lawyers and programmers to learn new concepts (logic programming, ontologies, defeasability, deontic logic, etc.) before they can put the tools to use. The risk is that these artifact meant as reusable and accessible models create a third class of model-makers, distinct from lawyers and programmers, that act as an intermediary and barrier that controls how the worlds of CS and Law interact with each other.

These developments strike a heavy blow to the dream of machine-consumable and infinitely reusable legal modelling that would be conform to the shared utopia of legal formalism and cybernetics. Consequently, the third decade of AI & Law retrospective is dominated by machine learning approaches that aim to directly solve a specific task without going through the intermediate step of making a general, reusable model of the law being analyzed first. However, the hopes have been recently revived by the latest developments in AI & Law's parent field, namely generative Large Language Models (LLMs) such as ChatGPT. Spearheading this research effort is Stanford CodeX's Doulcet [2023] that has tried to use ChatGPT to directly generate a formal and reusable model of the law from the legal text. However, he quickly realized that the reusable model he created that way was less efficient than using ChatGPT to translate directly legal texts to specific pieces of executable code instead: "*[...] once you can automatically translate any piece of legislation/contract into code, the most interoperable format is the words themselves. And as the cost of building a code representation of a legal text goes down, you care less about reusability of your legislation model for multiple use cases. [...] Instead we need to build systems that automatically transform legal text into some code in a popular programming language that solves the task at hand. Using popular programming language, it is easier to let the LLM generate the code (it likely already knows Python or Typescript), and we get better tooling!*". According to Doulcet's observations, technology may have gone full circle and in the end, the correct way of modelling the legal text is... the legal text itself. More pragmatically, tenants of logic programming and ontologies are currently trying to integrate LLMs into their modelling processes, in order to compensate the limits of their technology of choice and try and finally achieve the dream of machine-consumable and infinitely reusable legal modelling. But is this dream really desirable? Is it worth spending another decade of collective work trying to achieve it at the expense of other goals for the research at the intersection of CS and Law? Here, it is useful to recall the analysis by Leith [2016] about the fall of the movement of legal expert systems powered by logic programming in the 1990's: "*why was there optimism, was there ever any success, and–if as I suggest–there was none, then why was such a huge extravaganza of funding for expert systems research in a field (legal technology) which has been practically starved of funding in all the other decades outwith the 1980's? I tried to answer these questions in my* Formalism in AI and Computer Science *[Leith 1990], suggesting that the focus on the machine rather than the user had led technicians into fields which they little understood, and I still believe that was the underlying reason for the decade. [...] the AI community now rewrite their AI projects to suit funders who are less keen on the AI approach - XML technology being one such funding source*".

4 WANTING IT ALL IS A COUNTERPRODUCTIVE APPROACH

From now, I will try to generalize some arguments from the critique of the technical systems made in the previous section, about why the dream of creating machine-consumable and infinitely reusable legal models is counterproductive. These arguments are not philosophical or related to how the rule of law should be respected or not, but very much technically-oriented and come from the experience of engineering and applied science best practices. By providing these new ideas to the debate, I hope to reframe the current antagonism of bold disruptive technologist *vs.* conservative and scientifically illiterate lawyers with the more nuanced proposition that good engineering and legal technological work will benefit from detaching from the ideologically-imposed goals and methods discussed in the previous sections. First, I'll argue that a general-purpose legal model can never be precise enough to cover the precision needs of critical applications. Second, infinitely reusable legal modelling is an unscoped system that is impossible to evaluate correctly, thus evading any rigorous scientific process. Third, the quest for the all-encompassing model requires to put a lot of infrastructure work upfront before reaping the benefits, which also means wasting a lot of time if the approach end up not working.

My first argument builds on the notion of inferencing bias of Valente et al. [1999]: "*knowledge is usually modeled with certain types of inferences in mind. For example, if we expect to use the Loom classifier to infer whether or not two interval smeet (that is,* `(meets int1 int2)`*), we need to add enough information in the definition of the relation meets to enable the classifier to use it. If, however, we only want to assert that the intervals meet and use this information for other inferences, it is enough to state the range and domain of the meets relation*". Similarly, real-world applications of AI & Law need to worry about number representations, rounding errors, data structure layout, error handling, null data, etc. All of these actually change the result returned to the user, so the question becomes: what is the level of precision we expect from AI & Law tools? I would argue that the killer applications of AI & Law are also the ones where the stakes are high and we need computer precision to compensate for human deficiencies. But an all-encompassing, reusable model is more likely to spread out its precision and accuracy over its diverse uses. Minimizing the precision problem saying that the tool is only here to "assist" is not satisfactory in my opinion: a tool for which you have to commit a lot of infrastructure and time to get answers that are often wrong is not convivial in the sense of Illich and Lang [1973].

The second argument points out the impossibility to correctly evaluate the usefulness of performance of a general, infinitely-reusable model (of law or of something else). This is a broader problem for the field of AI. The evaluation problem in the field of AI & Law is correctly pointed by Novotná and Libal [2022]: "*In [Cohen and Howe 1988], the authors state that the evaluation of the experiments and the methods 'expedites the understanding of available methods and so their integration into further research'. The authors in [Conrad and Zeleznikow 2015] argue, that 'a performance-based ethic signifies a level of maturity and scientific rigor within a community'. However, the meta-analysis of research studies in the field of artificial intelligence and law in [Conrad and Zeleznikow 2013, 2015; Hall and Zeleznikow 2001] shows that great part of studies does not contain any kind of evaluation whatsoever*". Machine learning researchers have addressed the evaluation problem by building shared benchmarks and evaluation metrics like Imagenet [Deng et al. 2009]. These benchmarks have biases and may lead to overfitting models, but they are a necessary step for advancing the field. Works by Holzenberger et al. [2020] or Guha et al. [2023] are starting to fill that gap but the practice should extend to non-machine-learning-based AI & Law too. Otherwise, the field's quest will be similar to the unscoped quest for AGI described by Gebru and Torres [2023].

The third argument concerns the very prospective nature of making general models that we hope can be reused afterwards. AI & Law projects that start without having identified precisely who are the users and who are the developers of their tools are most likely doomed to miss any target they later set for themselves. The technical choices should follow the needs of the users and not vice-versa. By choosing architectures and objectives with only the dreams legal formalism and cybernetics in mind, the field will continue hitting hard barriers for adoptions by institutions and companies that operate under different ideologies and conceptions of legal practice. While the theoretical foundation of AI & Law have been deeply studied, the applied research branch of the field is missing killer applications and adoption. In applied research, it is not sufficient to "propose an innovative architecture" or "offer some ideas on whether certain techniques can help users"; the goal is to have technical solutions battle-tested and some of them ultimately adopted to become industry standards.

5 CONCLUSION: BRING BACK THE SCIENTIFIC METHOD

I personally come from a research field–formal methods–whose relationship with applied research and practical applications is difficult [MacKenzie 2004]. Formal methods, coincidentally, are also a scientific discipline born out of the early field of AI. But today, formal methods have dropped their affiliation to AI and have resolutely chosen concrete areas of applications: critical software, embedded systems, model checking, etc. Some tools from formal methods have become industry standards in railroads, avionics, the nuclear sector, etc. At the same time, the theoretical activity in formal methods is still significant, and we're starting too see thirty-years-old theoretical foundations (such as linear logic by Girard [1987]) being used as the basis for popular programming languages (such as Rust). What I retain from my experience in formal methods is the healthy divide between the evaluation criteria for theoretical *vs.* applied research, even when faced with contributions that are basically formalisms or models. If your contribution is theoretical, then you must show how it is more expressive, concise, elegant, etc. than related work. If your contribution is applied, then you must show how it is more performant, practical, adopted, etc. than related work. Negative results can be contributions but only if the paper explains why the results are negative and what we can learn from it (other than something does not work). In short, follow the scientific method! Deviating from it threatens the ability of the field to weed out unproductive approaches and allocates its research ressources optimally. And if funding bodies continue to be fascinated by legal formalism, cybernetics or the tantalizing promise of government by machines, then it is up to us, as peers in our scholarly venues, and not as cogs of a system caught up in a feedback loop, to judge the value of contributions, with technical and legal sharpness alike.

REFERENCES

Layman E Allen. 1956. Symbolic logic: A razor-edged tool for drafting and interpreting legal documents. *Yale LJ* 66 (1956), 833.

Joaquín Arias, Mar Moreno-Rebato, Jose A Rodriguez-García, and Sascha Ossowski. 2021. Modeling administrative discretion using goal-directed answer set programming. In *Conference of the Spanish Association for Artificial Intelligence*. Springer, 258–267.

Tara Athan, Harold Boley, Guido Governatori, Monica Palmirani, Adrian Paschke, and Adam Wyner. 2013. Oasis legalruleml. In *proceedings of the fourteenth international conference on artificial intelligence and law*. 3–12.

Tara Athan, Guido Governatori, Monica Palmirani, Adrian Paschke, and Adam Wyner. 2015. LegalRuleML: Design principles and foundations. In *Reasoning Web International Summer School*. Springer, 151–188.

Gioele Barabucci, Luca Cervone, Angelo Di Iorio, Monica Palmirani, Silvio Peroni, and Fabio Vitali. 2010. Managing semantics in XML vocabularies: an experience in the legal and legislative domain. In *Proceedings of Balisage: The markup conference*, Vol. 5.

Shrutarshi Basu, Nate Foster, and James Grimmelmann. 2019. Property conveyances as a programming language. In *Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. 128–142.

Kent Beck, Mike Beedle, Arie Van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler, James Grenning, Jim Highsmith, Andrew Hunt, Ron Jeffries, et al. 2001. Manifesto for agile software development.

Danièle Bourcier. 2017. Artificial intelligence and law: what perspective? Presentation at the Congress of the World Organisation of Systems and Cybernetics.

Joost Breuker, André Valente, and Radboud Winkels. 2004. Legal ontologies in knowledge engineering and information management. *Artificial intelligence and law* 12 (2004), 241–277.

Christopher Brewster and Kieron O'Hara. 2004. Knowledge representation with ontologies: the present and future. *IEEE Intelligent Systems* 19, 1 (2004), 72–81.

Pierre Catala, Lucien Mehl, and Edmond Bertrand. 1974. *Constitution et exploitation informatique d'un ensemble documentaire en droit (droit de l'urbanisme et de de la construction).* Ed. du CNRS.

Paul R Cohen and Adele E Howe. 1988. How evaluation guides AI research: The message still counts more than the medium. *AI magazine* 9, 4 (1988), 35–35.

Alain Colmerauer and Philippe Roussel. 1996. The birth of Prolog. In *History of programming languages—II*. 331–367.

Jack G. Conrad and John Zeleznikow. 2013. The Significance of Evaluation in AI and Law: A Case Study Re-Examining ICAIL Proceedings. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law* (Rome, Italy) *(ICAIL '13)*. Association for Computing Machinery, New York, NY, USA, 186–191. https://doi.org/10.1145/2514601.2514624

Jack G. Conrad and John Zeleznikow. 2015. The Role of Evaluation in AI and Law: An Examination of Its Different Forms in the AI and Law Journal. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law* (San Diego, California) *(ICAIL '15)*. Association for Computing Machinery, New York, NY, USA, 181–186. https://doi.org/10.1145/2746090.2746116

Giuseppe Contissa, Francesco Godano, and Giovanni Sartor. 2021. Computation, cybernetics and the law at the origins of legal informatics. *Italian Philosophy of Technology: Socio-Cultural, Legal, Scientific and Aesthetic Perspectives on Technology* (2021), 91–110.

Accord Contributors. 2018. The Accord Project. https://accordproject.org/.

Silvia Crafa, Cosimo Laneve, Giovanni Sartor, and Adele Veschetti. 2023. Pacta sunt servanda: legal contracts in Stipula. *Science of Computer Programming* 225 (2023), 102911.

Anthony D'Amato. 1977. Can/should computers replace judges? *Georgia Law Review* 11 (1977), 11–36.

Tim de Sousa and Pia Andrews. 2019. When we code the rules on which our society runs, we can create better results and new opportunities for the public and regulators, and companies looking to make compliance easier. *The Mandarin* (2019).

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.

John Dewey. 1924. Logical method and law. *Cornell LQ* 10 (1924), 17.

L. Diver, P. McBride, M. Medvedeva, A. Banerjee, E. D'hondt, T. Duarte, D. Dushi, G. Gori, E. van den Hoven, P. Meessen, and M. Hildebrandt. 2022. Typology of Legal Technologies. COHUBICOL. https://publications.cohubicol.com/typology

Pierre-Loïc Doulcet. 2023. Computational Law, AI status update. https://twitter.com/hexapode/status/1691386964170207232.

Vimal Dwivedi, Alex Norta, Alexander Wulf, Benjamin Leiding, Sandeep Saxena, and Chibuzor Udokwu. 2021. A formal specification smart-contract language for legally binding decentralized autonomous organizations. *IEEE access* 9 (2021), 76069–76082.

Lexon Foundation. 2019. Lexon. http://www.lexon.tech/.

Hamish Fraser. 2021. What is Better Rules? https://www.digital.govt.nz/blog/what-is-better-rules/.

Vittorio Frosini. 1968. *Cibernetica, Diritto e società*. Milano.

Vittorio Frosini. 1973. L'automazione elettronica nella giurisprudenza e nell'amministrazione pubblica. *Bollettino bibliografico d'informatica generale e applicata al diritto* 2, 3-4 (1973), 101–104.

Timnit Gebru and Émile Torres. 2023. Eugenics and the Promise of Utopia through Artificial General Intelligence. Keynote presentation at the SaTML conference. https://www.youtube.com/watch?v=P7XT4TWLzJw

Michael Genesereth. 2015. Computational Law: The Cop in the Backseat. Stanford CodeX Center.

Jean-Yves Girard. 1987. Linear logic. *Theoretical computer science* 50, 1 (1987), 1–101.

Guido Governatori. 2015. The Regorous approach to process compliance. In *2015 IEEE 19th International Enterprise Distributed Object Computing Workshop*. IEEE, 33–40.

Guido Governatori, Trevor Bench-Capon, Bart Verheij, Michał Araszkiewicz, Enrico Francesconi, and Matthias Grabmair. 2022. Thirty years of Artificial Intelligence and Law: the first decade. *Artificial Intelligence and Law* 30, 4 (01 Dec 2022), 481–519. https://doi.org/10.1007/s10506-022-09329-4

Neel Guha, Julian Nyarko, Daniel E. Ho, Christopher Ré, Adam Chilton, Aditya Narayana, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel N. Rockmore, Diego Zambrano, Dmitry Talisman, Enam Hoque, Faiz Surani, Frank Fagan, Galit Sarfaty, Gregory M. Dickinson, Haggai Porat, Jason Hegland, Jessica Wu, Joe Nudell, Joel Niklaus, John Nay, Jonathan H. Choi, Kevin Tobia, Margaret Hagan, Megan Ma, Michael Livermore, Nikon Rasumov-Rahe, Nils Holzenberger, Noam Kolt, Peter Henderson, Sean Rehaag, Sharad Goel, Shang Gao, Spencer Williams, Sunny Gandhi, Tom Zur, Varun Iyer, and Zehua Li. 2023. LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models. arXiv:2308.11462 [cs.CL]

Maria Jean J Hall and John Zeleznikow. 2001. Acknowledging insufficiency in the evaluation of legal knowledge-based systems: strategies towards a broadbased evaluation model. In *Proceedings of the 8th international conference on Artificial intelligence and law*. 147–156.

Xiao He, Bohan Qin, Yan Zhu, Xing Chen, and Yi Liu. 2018. Spesc: A specification language for smart contracts. In *2018 IEEE 42nd Annual computer software and applications conference (COMPSAC)*, Vol. 1. IEEE, 132–137.

Mireille Hildebrandt. 2020. Code-driven Law: Freezing the Future and Scaling the Past. In *Is Law Computable? : Critical Perspectives on Law and Artificial Intelligence*, Simon Deakin and Christopher Markou (Eds.). Hart Publishing, Oxford, 67–84.

Nils Holzenberger, Andrew Blair-Stanek, and Benjamin Van Durme. 2020. A Dataset for Statutory Reasoning in Tax Law Entailment and Question Answering. *arXiv preprint arXiv:2005.05257* (2020).

Ivan Illich and Anne Lang. 1973. Tools for conviviality. (1973).

Viktor Knapp and Vladimir Vrecion. 1970. The Possibility of applying cybernetic methods to law and administration. *International social science journal* XXII, 3 (1970). https://unesdoc.unesco.org/notice?id=p::usmarcdef_0000024220

Philipp Körner, Michael Leuschel, João Barbosa, Vítor Santos Costa, Verónica Dahl, Manuel V Hermenegildo, Jose F Morales, Jan Wielemaker, Daniel Diaz, Salvador Abreu, et al. 2022. Fifty years of Prolog and beyond. *Theory and Practice of Logic Programming* 22, 6 (2022), 776–858.

Ho-Pun Lam and Guido Governatori. 2009. The making of SPINdle. In *Rule Interchange and Applications: International Symposium, RuleML 2009, Las Vegas, Nevada, USA, November 5-7, 2009. Proceedings 3*. Springer, 315–322.

Philip Leith. 1990. *Formalism in AI and computer science*. Ellis Horwood.

Philip Leith. 2016. The rise and fall of the legal expert system. *International Review of Law, Computers & Technology* 30, 3 (2016), 94–106.

How Khang Lim, Avishkar Mahajar, Martin Strecker, and Meng Weng Wong. 2022. Automating defeasible reasoning in law with answer set programming. (2022).

Donald MacKenzie. 2004. *Mechanizing proof: computing, risk, and trust*. MIT Press.

Eden Medina. 2006. Designing freedom, regulating a nation: socialist cybernetics in Allende's Chile. *Journal of Latin American Studies* 38, 3 (2006), 571–606.

Masha Medvedeva, Martijn Wieling, and Michel Vols. 2023. Rethinking the field of automatic prediction of court decisions. *Artificial Intelligence and Law* 31, 1 (2023), 195–212.

Lucien Mehl. 1957. La cybernetique et l'administration. *La Revue administrative* Juillet-Août, 58 (1957).

Lucien Mehl. 1959. Automation in the legal world: from the machine processing of legal information to the "Law Machine". In *Mechanisation of Thought Processes: Proceedings of a Symposium Held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*. Vol. II, Her Majesty's Stationery Office, London. https://aitopics.org/download/classics:97D0F0CA

Lucien Mehl and Jean-Marie Breton. 1970. L'automatisation de la recherche de l'information juridique par le procédé Docilis : Documents, interrogations libres. (1970).

Denis Merigoux. 2021. *Proof-oriented domain-specific language design for high-assurance software*. Theses. Université Paris sciences et lettres. https://tel.archives-ouvertes.fr/tel-03622012

Denis Merigoux, Nicolas Chataing, and Jonathan Protzenko. 2021. Catala: A Programming Language for the Law. *Proc. ACM Program. Lang.* 5, ICFP, Article 77 (Aug. 2021), 29 pages. https://doi.org/10.1145/3473582

James Mohun and Alex Roberts. 2020. Cracking the code: Rulemaking for humans and machines. (2020).

Tereza Novotná and Tomer Libal. 2022. An evaluation of methodologies for legal formalization. In *International Workshop on Explainable, Transparent Autonomous Agents and Multi-Agent Systems*. Springer, 189–203.

OECD. 2019. *Embracing Innovation in Government: Global Trends 2019*. Technical Report. Observatory for Public Sector Innovation.

Anja Oskamp and Marc Lauritsen. 2002. AI in Law Practice? So far, not much. *Artificial Intelligence and Law* 10, 4 (01 Dec 2002), 227–236. https://doi.org/10.1023/A:1025402013007

M. Palmirani, L. Cervone, O. Bujor, and M. Chiappetta. 2013. RAWE: an editor for rule markup of legal texts. In *CEUR Workshop Proceedings, Seattle, USA, 11–13 July 2013*, P. Fodor, D. Roman, D. Anicic, A. Wyner, M. Palmirani, D. Sottara, and F. Lévy (Eds.).

Monica Palmirani, Giuseppe Contissa, and Rossella Rubino. 2009. Fill the gap in the legal knowledge modelling. In *Rule Interchange and Applications: International Symposium, RuleML 2009, Las Vegas, Nevada, USA, November 5-7, 2009. Proceedings 3*. Springer, 305–314.

Monica Palmirani and Guido Governatori. 2018. Modelling Legal Knowledge for GDPR Compliance Checking. In *JURIX*, Vol. 313. 101–110.

Monica Palmirani, Guido Governatori, Antonino Rotolo, Said Tabet, Harold Boley, and Adrian Paschke. 2011. LegalRuleML: XML-based rules and norms. In *International Workshop on Rules and Rule Markup Languages for the Semantic Web*. Springer, 298–312.

Monica Palmirani, Michele Martoni, Arianna Rossi, Cesare Bartolini, and Livio Robaldo. 2018. PrOnto: privacy ontology for legal reasoning. In *International Conference on Electronic Government and the Information Systems Perspective*. Springer, 139–152.

Monica Palmirani and Fabio Vitali. 2011. Akoma-Ntoso for legal documents. *Legislative XML for the Semantic Web: Principles, Models, Standards for Document Management* (2011), 75–100.

Maurizio Parton, Marco Angelone, Carlo Metta, Stefania D'Ovidio, Roberta Massarelli, Luca Moscardelli, Gianluca Amato, and Cristiano De Nobili. 2024. Artificial intelligence and renegotiation of commercial lease contracts affected by pandemic-related contingencies from covid-19. The project A.I.A.Co. *Journal of Cross-disciplinary Research in Computational Law* 1, 4 (2024). https://www.youtube.com/watch?v=8AisFKcrJzY (forthcoming).

Silvio Peroni, Monica Palmirani, and Fabio Vitali. 2017. UNDO: The United Nations system document ontology. In *The Semantic Web–ISWC 2017: 16th International Semantic Web Conference, Vienna, Austria, October 21-25, 2017, Proceedings, Part II 16*. Springer, 175–183.

Ursula Plesner and Lise Justesen. 2022. The Double Darkness of Digitalization: Shaping Digital-ready Legislation to Reshape the Conditions for Public-sector Digitalization. *Science, Technology, & Human Values* 47, 1 (2022), 146–173. https://doi.org/10.1177/0162243921999715 arXiv:https://doi.org/10.1177/0162243921999715

Joseph Potvin. 2023. *Data with direction: design research leading to a system specification for 'an internet of rules'*. Ph. D. Dissertation. Université du Québec en Outaouais.

Joseph Potvin, Don Kelly, William Olders, Wayne Cunneyworth, Ryan Fleck, Craig Atkinson, Calvin Hutcheon, Ted Kim, Angela Bernal, and Stéphane Gagnon. 2021. Oughtomation: Practical Normative Data. (2021). https://gitlab.com/xalgorithms-alliance/oughtomation-paper/-/raw/master/ThesisJPotvin_Oughtomation_EXCERPT_1-87+references_draft-2021-05-09PDF.pdf

Eric S Raymond. 2003. *The art of Unix programming*. Addison-Wesley Professional.

Edwina L Rissland, Kevin D Ashley, and Ronald Prescott Loui. 2003. AI and Law: A fruitful synergy. *Artificial Intelligence* 150, 1-2 (2003), 1–15.

Giovanni Sartor, Michał Araszkiewicz, Katie Atkinson, Floris Bex, Tom van Engers, Enrico Francesconi, Henry Prakken, Giovanni Sileno, Frank Schilder, Adam Wyner, and Trevor Bench-Capon. 2022. Thirty years of Artificial Intelligence and Law: the second decade. *Artificial Intelligence and Law* 30, 4 (01 Dec 2022), 521–557. https://doi.org/10.1007/s10506-022-09326-7

James C. Scott. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press. http://www.jstor.org/stable/j.ctt1nq3vk

M. J. Sergot, A. S. Kamble, and K. K. Bajaj. 1991. Indian Central Civil Service Pension Rules: A Case Study in Logic Programming Applied to Regulations. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law* (Oxford, England) *(ICAIL '91)*. Association for Computing Machinery, New York, NY, USA, 118–127. https://doi.org/10.1145/112646.112661

M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, and H. T. Cory. 1986. The British Nationality Act As a Logic Program. *Commun. ACM* 29, 5 (May 1986), 370–386.

Giovanni Sileno. 2016. *Aligning Law and Action*. Ph. D. Dissertation. University of Amsterdam. https://gsileno.net/articles/sileno.PhDthesis.pdf

Didier Torny. 2005. L'administration sanitaire entre contraintes techniques et contraintes juridiques. L'exemple des maladies émergentes. *Revue générale de droit médical* 6 (2005). https://halshs.archives-ouvertes.fr/halshs-01249084

Andre Valente, Thomas Russ, Robert MacGregor, and William Swartout. 1999. Building and (re) using an ontology of air campaign planning. *IEEE Intelligent Systems and their Applications* 14, 1 (1999), 27–36.

Serena Villata, Michal Araszkiewicz, Kevin Ashley, Trevor Bench-Capon, L. Karl Branting, Jack G. Conrad, and Adam Wyner. 2022. Thirty years of artificial intelligence and law: the third decade. *Artificial Intelligence and Law* 30, 4 (01 Dec 2022), 561–591. https://doi.org/10.1007/s10506-022-09327-6

Joseph Weizenbaum. 1976. Computer power and human reason: From judgment to calculation. (1976).

Jan Wielemaker, Tom Schrijvers, Markus Triska, and Torbjörn Lager. 2012. Swi-prolog. *Theory and Practice of Logic Programming* 12, 1-2 (2012), 67–96.

A. Wright, D. Roon, and ConsenSys AG. 2019. Open Law. https://www.openlaw.io/.