

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Introduction: Legal and Ethical Dimensions of AI, NorMAS, and the Web of Data

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Ugo Pagallo, M.P. (2018). Introduction: Legal and Ethical Dimensions of AI, NorMAS, and the Web of Data. Berlin : Springer [10.1007/978-3-030-00178-0_1].

Availability:

This version is available at: <https://hdl.handle.net/11585/665649> since: 2019-02-14

Published:

DOI: http://doi.org/10.1007/978-3-030-00178-0_1

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Pagallo U., Palmirani M., Casanovas P., Sartor G., Villata S. (2018)

Introduction: Legal and Ethical Dimensions of AI, NorMAS, and the Web of Data

In: Pagallo U., Palmirani M., Casanovas P., Sartor G., Villata S. (eds) *AI Approaches to the Complexity of Legal Systems*.

Lecture Notes in Computer Science, vol 10791.

Springer, Cham.

First Online 23 October 2018

Print ISBN 978-3-030-00177-3

Online ISBN 978-3-030-00178-0

The final published version is available online at:

https://doi.org/10.1007/978-3-030-00178-0_1

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

<https://www.springer.com/gp/open-access/publication-policies/self-archiving-policy>

Introduction: Legal and Ethical Dimensions of AI, NorMAS, and the Web of Data

Ugo Pagallo

ugo.pagallo@unito.it

Torino Law School, University of Torino, Lungo Dora Siena 100, 10153 Turin, Italy

Monica Palmirani

monica.palmirani@unibo.it

CIRSFID, University of Bologna, via Zamboni 33, 40126 Bologna, Italy

Giovanni Sartor

giovanni.sartor@unibo.it

CIRSFID, University of Bologna, via Zamboni 33, 40126 Bologna, Italy

Pompeu Casanovas

pompeu.casanovas@uab.cat

CRC D2D, La Trobe Law School, Bundoora Campus, Melbourne, Australia ; IDT, Autonomous University of Barcelona, Bellaterra, Spain

Abstract. AICOL workshops aim to bridge the multiple ways of understanding legal systems and legal reasoning in the field of AI and Law. Moreover, they pay special attention to the complexity of both legal systems and legal studies, on one hand, and the expanding power of the internet and engineering applications, on the other. Along with a fruitful interaction and exchange of methodologies and knowledge between some of the most relevant contributions to AI work on contemporary legal systems, the goal is to integrate such a discussion with legal theory, political philosophy, and empirical legal approaches. More particularly, we focus on four subjects, namely, (i) language and complex systems in law; (ii) ontologies and the representation of legal knowledge; (iii) argumentation and logics; (iv) dialogue and legal multimedia.

Keywords: AI & law Legal theory Complex systems Semantic web Legal ontologies Legal semantic web services Argumentation

1 Introduction

The first volume of this workshop series on Artificial Intelligence approaches to the complexity of legal systems (AICOL) was released at the very beginning of this decade (2010). In the meanwhile, the field of Artificial Intelligence ("AI") has known a new renaissance: for instance, according to the tally Google provided to MIT Technology Review in March 2017, the company published 218 journal or conference papers on machine learning in 2016 alone, nearly twice as many as it did two years before [1]. Google's AI explosion illustrates a more general trend that has to do with the improvement of more sophisticated statistical and probabilistic methods, the increasing availability of large amount of data and of cheap, enormous computational power, up to the transformation of places and spaces into AI-friendly environments, e.g. smart cities

and domotics. All these factors have propelled a new ‘summer’ for AI. After the military and business sectors, AI applications have entered into people’s lives. From getting insurance to landing credit, from going to college to finding a job, even down to the use of GPS for navigation and interaction with the voice recognition features on our smartphones, AI is transforming and reshaping people’s daily interaction with others and their environment. Whereas AI apps and systems often go hand-in-hand with the breath-taking advancements in the field of robotics, the internet of things, and more, we can grasp what is going on in this field in different ways [2]. Suffice it to mention in this context two of them. First, according to the Director of the Information Innovation Office (I2O) at the Defense Advanced Research Projects Agency (DARPA) in the U.S. Department of Defense, John Launchbury, there would have been so far two waves of research in AI.¹

The first wave concerns systems based on “handcrafted knowledge,” such as programs for logistics scheduling, programs that play chess, and in the legal domain, TurboTax. Here, experts turn the complexity of the law into certain rules, and “the computer then is able to work through these rules.” Although such AI systems excel at complex reasoning, they were inadequate at perception and learning. In order to overcome these limits, we thus had to wait for the second wave of AI, that is, systems based on “statistical learning” and the “manifold hypothesis.” As shown by systems for voice recognition and face recognition, the overall idea is that “natural data forms lower- dimensional structures (manifolds) in the embedding space” and that the task for a learning system is to separate these manifolds by “stretching” and “squashing” the space.²

Along these lines, Richard and Daniel Susskind similarly propose to distinguish between a first generation and a second generation of AI systems, namely, between expert systems technologies and systems characterized by major progress in Big Data and search [3]. What this new wave of AI entails has to do more with the impact of AI on society and the law, than the law as a rich test bed and important application field for logic-based AI research. Whether or not the first wave aimed to replicate knowledge and reasoning processes that underpin human intelligence as a form of deductive logic, the second wave of AI brings about “two possible futures for the professions,” that is, either a more efficient version of the current state of affairs, or a profound transformation that will displace much of the work of traditional professionals. According to this stance, we can thus expect that “in the short and medium terms, these two futures will be realized in parallel.” Yet, the thesis of Richard and Daniel Susskind is that the second future will prevail: in their phrasing, “we will find new and better ways to share expertise in society, and our professions will steadily be dismantled” [3].

Against this highly problematical backdrop, three different levels of analysis should be however differentiated. They concern the normative challenges of the second wave of AI from a political, theoretical, and technical viewpoint, namely (i) the political decisions that should be—or have already been—taken vis-à-vis current developments

¹ See the video entitled “A DARPA Perspective on Artificial Intelligence,” available online at <https://youtube/-O01G3tSYpU>.

² Ibid.

of e.g. self-driving cars, or autonomous lethal weapons; (ii) the profound transformations that affect today's legal systems vis-à-vis the employment of e.g. machine learning techniques; and, (iii) the advancements in the state-of-the-art that regard such areas, as semantic web applications and language knowledge management in the legal domain, or ejustice advanced applications. Each one of these levels of analysis is deepening in the following sections.

2 Architectural Challenges

The political stance on current developments of AI hinges on a basic fact: the more the second wave of AI advances, the more AI impacts on current pillars of society and the law, so that political decisions will have to be taken as regards some AI applications, such as lethal autonomous weapons, or self-driving cars. Over the past years, scholars, non-profit organizations, and institutions alike have increasingly stressed the ethical concerns and normative challenges brought about by many autonomous and intelligent system designs [4–6]. The aim of the law to govern this field of technological innovation suggests that we should distinguish between two different levels of political intervention, that is, either through the primary rules of the law, or through its secondary rules [7].

According to the primary rules of the law, the goal is to directly govern social and individual behaviour through the menace of legal sanctions. Legislators have so far aimed to attain this end through methods of accident control that either cut back on the scale of the activity via, e.g., strict liability rules, or intend to prevent such activities through bans, or the precautionary principle. Regulations can be divided into four different categories, that is, (a) the regulation of human producers and designers of AI systems through law, e.g. either through ISO standards or liability norms for users of AI; (b) the regulation of user behaviour through the design of AI, that is, by designing AI systems in such a way that unlawful actions of humans are not allowed; (c) the regulation of the legal effects of AI behaviour through the norms set up by lawmakers, e.g. the effects of contracts and negotiations through AI applications; and, (d) the regulation of AI behaviour through design, that is, by embedding normative constraints into the design of the AI system [8].

Current default norms of legal responsibility can entail however a vicious circle, since e.g. strict liability rules—let aside bans, or the precautionary principle—may end up hindering research and development in this field. The recent wave of extremely detailed regulations on the use of drones by the Italian Civil Aviation Authority, i.e. “ENAC,” illustrates this deadlock. The paradox stressed in the field of web security decades ago, could indeed be extended with a pinch of salt to the Italian regulation on the use of drones as well: the only legal drone would be “one that is powered off, cast in a block of concrete and sealed in a lead-lined room with armed guards – and even then I have my doubts.” [9] As a result, we often lack enough data on the probability of events, their consequences and costs, to determine the levels of risk and thus, the amount of insurance premiums and further mechanisms, on which new forms of accountability for the behaviour of such systems may hinge. How, then, can we prevent legislations that may hinder the research in AI? How should we deal with the peculiar

unpredictability and risky behaviour of some AI systems? How should we legally regulate the future?

A feasible way out can be given by the secondary rules of the law, namely, the rules of the law that create, modify, or suppress the primary rules of the system. Among the multiple legal techniques with which we can properly address the normative challenges of the second wave of AI, suffice it to mention here three of them. First, focus should be on Justice Brandeis's doctrine of experimental federalism, as espoused in *New State Ice Co. v. Leibmann* (285 US 262 (1932)). The idea is to flesh out the content of the rules that shall govern individual behaviour through a beneficial competition among legal systems. This is what occurs nowadays in the field of self-driving cars in the US, where several states have enacted their own laws for this kind of technology. At its best possible light, the same policy will be at work with the EU regulation in the field of data protection [10, 11].

Second, attention should be drawn to the principle of implementation neutrality, according to which regulations are by definition specific to that technology and yet do not favour one or more of its possible implementations. The 2016 Federal Automated Vehicles Policy of the U.S. Department of Transportation illustrates this legal technique. Although regulations are by definition specific to that technology, e.g. autonomous vehicles, there is no favouritism for one or more of its possible implementations. Even when the law sets up a particular attribute of that technology, lawmakers can draft the legal requirement in such a way that non-compliant implementations can be modified to become compliant.

Third, legislators can adopt forms of legal experimentation. For example, over the past decade and a half, the Japanese government has worked out a way to address the normative challenges of robotics through the creation of special zones for their empirical testing and development, namely, a form of living lab, or *Tokku* [12]. Likewise, in the field of autonomous vehicles, several EU countries have endorsed this kind of approach: Sweden has sponsored the world's first large-scale autonomous driving pilot project, in which self-driving cars use public roads in everyday driving conditions; Germany has allowed a number of tests with various levels of automation on highways, e.g. Audi's tests with an autonomously driving car on highway A9 between Ingolstadt and Nuremberg.

In general terms, these forms of experimentation through lawfully de-regulated special zones represent the legal basis on which to collect empirical data and sufficient knowledge to make rational decisions for a number of critical issues. We can improve our understanding of how AI systems may react in various contexts and satisfy human needs. We can better appreciate risks and threats brought on by possible losses of control of AI systems, so as to keep them in check. We can further develop theoretical frameworks that allow us to better appreciate the space of potential systems that avoid undesirable behaviours. In addition, we can rationally address the legal aspects of this experimentation, covering many potential issues raised by the next-generation AI systems and managing such requirements, which often represent a formidable obstacle for this kind of research, as public authorizations for security reasons, formal consent for the processing and use of personal data, mechanisms of distributing risks through insurance models and authentication systems, and more. The different legal techniques and types of rules that lawmakers may employ, on the one hand, should not overlook

the importance of the goals and values that are at stake with choices of technological dependence, delegation and trust, in order to determine the good mix between legal automation and public deliberation [13]. On the other hand, such choices of technological dependence, delegation and trust, through AI systems and procedures of legal automation are affecting pillars and tenets of today's law. As stressed above in this introduction, AI technology profoundly affects both the requirements and functions of the law, namely, what the law is supposed to be (requirements), and what it is called to do (functions). This profound transformation has to be examined separately in the next section.

3 Ethical and Legal Challenges: Device and Linked Democracy

In one of the most celebrated 2014 John Klossner's cartoons on the Internet of Things the husband resignedly says to his wife: "We have to go out for dinner. The refrigerator isn't speaking to the stove."³ This is not a joke anymore, and neither is the possibility of connecting thousands of billions of devices that can literally speak to each other.

A world of smart objects shreds new challenges into the interconnected world of humans and machines. The 2015 IBM Institute for Business Value Report [14] has pointed out five major challenges: (i) the cost of connectivity (prohibitively high), (ii) the Internet after trust (in the after-Snowden era "trust is over" and "IoT solutions built as centralized systems with trusted partners is now something of a fantasy"), (iii) not-future proof (many companies are quick to enter the market but it is very hard to exit: the cost of software updates and fixes in products long obsolete and discontinued), (iv) a lack of functional value (lack of meaningful value creation), (v) broken business models (in information markets, the marginal cost of additional capacity—advertising—or incremental supply—user data—is zero).

This is setting the conditions for "Device Democracy", in which "devices are empowered to autonomously execute digital contracts such as agreements, payments and barter with peer devices by searching for their own software updates, verifying trustworthiness with peers, and paying for and exchanging resources and services. This allows them to function as self-maintaining, self-servicing devices" [14].

IBM suggests three new methodological trends for a scalable, secure, and efficient IoT regarding: (i) architecture (private-by-design), (ii) business and economic insights (key vectors of disruption), (iii) and product and user experience design (the transformation of physical products into meaningful digital experiences).

The keyword here is the emergence of "meaningful experiences" in between relationships, properties, and objects. Interestingly this has been also enhanced by 2016 and 2017 Gartner Hype Cycle of Emerging Technologies: (i) transparently immersive experiences (such as human augmentation), (ii) perceptual smart machines (such as

³ <https://www.computerworld.com/article/2858429/enterprise-applications/2014-the-tech-year-in-cartoons.html#slide5>.

personal analytics), and (iii) digital platforms (including blockchain technologies).⁴ “AI everywhere” and blockchain are especially highlighted.

Our contention is that meaning is created, distributed, and framed through a complex and dynamic world in which ethics and law cannot be set apart or let alone to regulate processes, actions, and outcomes. Thus, regulations are entrenched and evolve dynamically according to the network and the specific ecosystems they are contributing to create. Meaningful experiences and smart devices entail smart regulations.

This is a feature already stressed for many years now by all attempts to frame new developments in norMAS [15], non-standard deontic logic [16] and law and the semantic web [17]. What is new is the attention brought to encompass innovation, semantic developments and blockchain technologies with ethical, democratic (political), and legal values alike.

In a recent Nexxus Whitepaper,⁵ Gavin Andresen, the leader of Blockchain Foundation, equates crypto-currencies with empowering people: “an unstoppable grassroots movement that won’t be trampled on by any government or bank. It’s all about the people’s freedom and reclaiming it” [18].

The first generation produced blockchain-based network protocols, such as Bitcoin, Ethereum, Litecoin, and Zerocash. The second one can take this same idea of creating encrypted blocks further: it uses distributed ledger databases along with user-programmable smart contracts⁶ to create a number of social contracts in many fields—universal basic income schemes, birth and death certificates, business licenses, property titles, educational qualifications, marriage etc. This actually is the original dream of the semantic web. MIT Digital Currency Initiative furnishes good examples in the public domain, as they have launched the identity based services offered by BitNation, a project aimed at decentralising governance at a global scale, e.g. a World Citizenship ID based on blockchain, and a Refugee Emergency Response project. Ideally, when applied to social or economic institutions, the result would be Decentralized Autonomous Organizations (DAO): self-running organisations in which users can become part of the chain performing things that computers cannot do [19].⁷

However, there is still a long road ahead before implementing DAO properly, as complexity does not only lie on regulations, and transactions fuel a mixed, hybrid social and economic reality that displays its own problems and conflicts. The recent Ethereum crisis has shown security vulnerabilities [20]. Although preventive mechanisms such as distributed consensus, cryptography, and anonymity are put in place, blockchain technologies remain vulnerable to many types of risks—mainly: the 51% attack, account takeover, digital identity theft, money laundering, and hacking [21].

⁴ <https://www.gartner.com/smarterwithgartner/>.

⁵ Nexxus Partners was established in January 2016 in Texas, USA as a services company for the bitcoin and cryptocurrency industry.

⁶ “Smart contracts are computer programs that can be correctly executed by a network of mutually distrusting nodes, without the need of an external trusted authority” [20].

⁷ See the first Decentralized Autonomous Organization (DAO) code to automate organizational governance and decision-making at [19].

There are problems regarding the legal dimension as well. Some criticisms have already stated that “cryptocurrencies cannot solve the problem of incomplete [relational] contracts, and as long as contracts are incomplete, humans will need to resolve ambiguities” [22]. The same diffidence has been shown from a public legal standpoint, as crypto-currencies cannot build up by their own a new public space [23].

The other way around, there are economic interpretations that highlight its positive aspects. Ethereum and blockchain platforms have been received “as a new type of economy: a ‘spontaneous organization’, which is a self-governing organization with the coordination properties of a market (Hayek), the governance properties of a commons (Ostrom), and the constitutional properties of a nation state (Brenan and Buchanan)” [24].

Perhaps this syncretic view is too over-confident, equating different dimensions (economic, social, and legal) but, be as it may, law and the definition of law—what it counts for—are at stake: “the legal status of DAOs remains the subject of active and vigorous debate and discussion. Not everyone shares the same definition. [...] Ultimately, how a DAO functions and its legal status will depend on many factors, including how DAO code is used, where it is used, and who uses it.” [19]

It is our contention that the synergy between different kinds of complementary technologies can help to solve these regulatory puzzles and tensions, i.e. Blockchain is the result of assembling two software paradigms (peer-to-peer applications and distributed hash tables). They are not the only ones. Semantic technologies and linked data can be used to ease the tensions and create the shared scenarios in which crypto-currencies and smart contracts can be safely and effectively used in a personalised manner by a vast plurality of users. “Similarly as block chain technology can facilitate distributed currency, trust and contracts application, Linked Data facilitated distributed data management without central authorities” [25].

But for this to happen, to cross jurisdictions and different types of legal obstacles, smart regulations and values are essential and should be similarly linked and harmonised. Traditional legal tools at national, European and international levels, are important, but they still fall short to cope with the complexity of algorithm governance to reach metadata regulatory dimensions and layers [26]. This is why law, governance, and ethics are at the same time being embedded into design, and re-enacted again as contextually-driven to shape sustainable regulatory ecosystems. Beyond epistemic and deliberative democracy, one of the concepts that have recently coined to describe this new situation is linked democracy, i.e. the endorsement of (embedded) democratic values to preserve rights and protect people on the web of data [27]. It is worth noticing that these common trends are related to the combination of political crowdsourcing, legal and ethical argumentation, and expert knowledge [28]. Innovation is deemed to be a crucial component of democracy [29]. Thus, what the law is supposed to be and what it is called to do are related not only to its architecture and tools (e.g. normative systems, laws and rights) but to the many ways of balancing citizens’ compliance and participation.

4 Web of Data and Legal Analytics

Over the last two decades we have witnessed a remarkable volume of legal documents and legal big data being put out in open format (e.g., the legal XML movement). The information was represented using specific technical standards capable of modelling legal knowledge, norms, and concepts [17, 30] in machine-readable format.

NormInRete [31] is an XML standard the Italian government issued in 2001 as the official XML vocabulary for the country's legislative documents. MetaLex was created in 2002 in the Netherlands, and it evolved into CEN-MetaLex as a general format for the interoperability of legal documents across Europe, this thanks to the EU Project ESTRELLA [32]. Another significant outcome of the ESTRELLA project was the Legal Knowledge Interchange Format-LKIF, composed of two main pillars: (i) a core legal ontology [33] and (ii) a legal-rule language [34]. Even if these outcomes are encouraging, they lack a common-framework technical design making it possible to easily integrate all the Semantic Web layers (e.g., text, norms, ontology). For this reason, the Akoma Ntoso project [35] (an UNDESA-based African initiative)⁸ took the best practices from those experiences and in 2006 designed a unique XSD schema for all legal documents (e.g., including caselaw and UN resolutions [39]) and lawmaking traditions (e.g., common law and civil law). In 2012, LegalDocML TC,⁹ of OASIS, expanded the Akoma Ntoso XML vocabulary to embrace an international vision of legal-document annotation. OASIS's LegalCiteM¹⁰ TC provides semantic representation of legal references so as to foster a convergence of many existing syntaxes for legal and legislative identifiers, including ELI [37], ECLI [36], URN-LEX,¹¹ and the Akoma Ntoso Naming Convention [40], making sure that legal document collections can unambiguously be referred to and are also connectable to Linked Data assertions. OASIS's LegalRuleML TC [38] provides a standard for modelling constitutive and prescriptive norms using formal language for rules. LegalDocML, LegalRuleML, and LegalCiteM provide a common framework for modelling legal documents and for fostering contextual metadata. The CLOUD4EU [41, 42] project offers a rare example of a platform where those standards can act in an integrated manner: it is designed for the General Data Protection Regulation (GDPR), making it possible to provide compliance reports for this regulation.

LegalRuleML also provides an RDFS meta-model for modelling the deontic and defeasible logic operators applied in the legal domain in order to export metadata in RDF format. LegalDocML makes it possible to extract legal metadata and to convert it into RDF. In the web of data paradigm, RDF triples produce a distributed and networked legal knowledge repository that can be useful in enhancing the searchability of relevant legal concepts, the semantic classification of documents, a light legal-reasoning approach, and the integration of metadata with other nonlegal sources (e.g.,

⁸ United Nations Department of Economic and Social Affairs <https://www.un.org/development/desa/en/>.

⁹ <https://www.oasis-open.org/committees/legaldocml/>.

¹⁰ https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=legalcitem.

¹¹ <https://datatracker.ietf.org/doc/draft-spinosa-urn-lex/>.

DBpedia). However, the RDF technique is based on a collection of triples (assertions composed of a subject, a predicate, and an object), often not validated, and the underlying theory of inferential logic only approximates reality, like a map. In the complex nonmonotonic legal system, RDF reasoning is quite risky, especially in situations where antecedents are rebutted (e.g., retroactivity, modifications of modifications, exceptions to exceptions, suspensions of the application of norms, annulment, etc.). For this reason, the legal XML community is inclined to extract RDF assertions as a subproduct of legal XML documents that are authentically, officially validated by experts (e.g., principle of self-contained and self-explained assertions in a unique XML file validated by the expert). In the meantime, this large web of legal data connected with official digital legal documents (often available from official-journal portals) in open format provides a rare opportunity to apply artificial intelligence (AI) techniques in order to rethink the legal theory and their institutions. In particular, legal analytics (LA)—combining data science, artificial intelligence, machine learning (ML), natural language processing techniques, and statistical methods—can reuse that vast and varied body of legal big data, even if approximate, to infer new patterns and legal knowledge, so as to then predict new models and bring out hidden correlations, often providing unexpected insights into the relation between legal phenomena [43].

The automated application of learning methods to vast sets of examples makes it possible to reproduce human behaviour and improve upon it, exploiting their computational power to learn from successes and failures and thus improve performance. Thanks to these new techniques and the sheer volume of the legal web of data now available on the Internet, artificial intelligence has been able to advance from mock examples to a host of real-life applications: conceptual retrieval and ranking, speech and image recognition, question-answering, recommendations, translation, planning, autonomous mobile robots, etc. Machine learning has been extensively applied to text analytics, which refers to the use of multiple technologies—such as linguistic, statistical, and machine learning techniques—to capture the information content of textual sources. Considering that legal sources of all kinds are recorded in textual form (this is true of statutes, regulations, judicial and other types of decisions and opinions, legal doctrines, contracts), the application of text analytics to the law has huge potentials. And indeed a number of applications are emerging, making it possible to automatically classify and extract documents, identify principles in judicial decisions, and predict the outcome of judicial cases.

Legal analytics (LA) has the potential to contribute to different aspects of legal scholarship and practice. LA can advance legal informatics, making for a vast range of successful new legal applications in the public and private sectors alike. In particular, it supports the provision of AI applications, but overcoming the knowledge-acquisition bottleneck. LA can support legal research, in particular, through its ability to unveil hidden patterns in data and documents, revealing unseen features of the structure and the functioning of legal institutions. The insights provided by LA, in combination with computable models of the law, support the development of a new, empirically based understanding of the law. LA can improve the efficiency and effectiveness of legal institutions, while providing all legal actors with better knowledge of the law and of its application. However, the use of big-data analytics has also raised a number of ethical

issues that are already being discussed in several domains (targeted deceptive commercial and noncommercial communication, discrimination, manipulation of public opinion, etc.). Also already emerging in the law are some questionable practices, particularly where law enforcement tries to predict illegal or otherwise unwanted behaviour (e.g., a tendency to offend or reoffend). More generally, the knowledge provided by analytics brings new ways to assess, influence, and control behaviour, and these aspects have yet to be fully analysed. For instance, no study exists so far on how the ability to predict court decisions, even those of specific judges, could influence judicial decision-making.

The combination of LA and legal XML techniques improve our interpretation of the AI inferences. XML nodes provide structural information and contextual metadata that, in combination with the predictive assertions of LA, could be used to mitigate two important negative side effects of LA techniques: (i) the introduction of bias from the past experiences and mistakes (e.g., negative case-law, bad legal drafting practices in legislation, influences due to socio-historical conditions) that can reinforce a tendency to reiterate incorrect models and may impair the ability to creatively find brilliant new solutions in the future (e.g., through filter-bubble effects); (ii) the fragmentation of legal knowledge into separate sentences or isolated data without a logical connection making it possible to achieve a consistent legal and logical narrative flow (e.g., contextless prediction). XML nodes could provide the skeleton needed to reassemble the huge amount of unexpected insights produced by the LA layer.

Finally, also crucial is the usability and the easy access to legal knowledge produced by LA. It is essential that the outcome of LA and legal XML sources in web applications and new devices (e.g., augmented reality) be also understandable by people who are not legal experts, without reframing the message; and, at any event, it is also essential to provide clear mechanisms for explaining the algorithm decision-making process and outcome. In this effort to achieve transparent communication we can turn to human-computer interaction techniques, making it possible to create a fair environment in which to better communicate the legal concepts and principles extracted by LA. The legal design community is working to create new design patterns, looking to provide better ways of displaying content, in such a way that the legal community and end users (e.g., citizens) can place greater trust in LA and legal XML [44–46].

5 On the Content of This Volume

This new AICOL volume is divided into six parts. They concern (i) legal philosophy, conceptual analysis, and epistemic approaches; (ii) rules and norms analysis and representation; (iii) legal vocabularies and natural language processing; (iv) legal ontologies and semantic annotation; (v) legal argumentation; and, (vi) courts, adjudication and dispute resolution.

5.1 Legal Philosophy, Conceptual Analysis, and Epistemic Approaches

In the first part of this volume, four papers deal with matters of legal philosophy, conceptual analysis, and epistemic approaches. In *RoboPrivacy and the Law as “Meta-technology”*, Ugo Pagallo examines how a particular class of robotic applications, i.e. service robots, or consumer robots, may affect current legal frameworks of privacy and data protection. Instead of a one-way movement of social evolution from technology to law, a key component of the analysis concerns the aim of the law to govern technological innovation as well as human and artificial behaviour through the regulatory tools of technology. By distinguishing between the primary rules of the law and its secondary rules, e.g. forms of legal experimentation, the chapter illustrates some of the ways in which the secondary rules of the law may allow us to understand what kind of primary rules we may want for our robots. In *Revisiting Constitutive Rules*, Giovanni Sileno, Alexander Boer and Tom Van Engers investigate how behaviour relates to norms, i.e. how a certain conduct acquires meaning in institutional terms. By addressing the double function of the ‘count-as’ relation, generally associated to constitutive rules and mostly accounted for its classificatory functions, the chapter reconsiders the relation between constitutive rules and regulative rules, and introduces a preliminary account on the ontological status of constitution. In *The Truth in Law and Its Explication*, Hajime Yoshino discusses what types of truth play their role in law and in what way such types of truth can be explicated. Whereas the concept of truth in law is classified into three types of truth, i.e. truth as fact, truth as validity and truth as justice, the chapter provides their formal semantic foundation, and analyses both the ways to explicate truth as validity and truth as justice in terms of intensional and extensional explication, and how to grasp the reasoning of justification and of creation. In *From Words to Images Through Legal Visualization*, Arianna Rossi and Monica Palmirani discuss the process of sense-making and interpretation of visual legal concepts that have been introduced in legal documents to make their meaning clearer and more intelligible. Whilst visualizations have also been automatically generated from semantically-enriched legal data, the analysis of current approaches to this subject represents the starting point to propose an empirical methodology that is inspired by the interaction with design practices and that will be tested in the future stages of the research.

5.2 Rules and Norms Analysis and Representation

The second part of the volume, which has to do with rules and norms analysis and representation, comprises seven contributions. In *A Petri Net-based Notation for Normative Modelling: Evaluation on Deontic Paradoxes*, Giovanni Sileno, Alexander Boer and Tom Van Engers focus on some of the problems that derive from the development of systems operating in alignment with norms, e.g. the continuous flow of events that modifies the normative directives under scrutiny. The chapter presents an alternative approach to some of these problems, by extending the Petri net notation to Logic Programming Petri Nets, so that the resulting visual formalism represents in an integrated, yet distinct fashion, procedural and declarative aspects of the system. In *Legal Patterns for Different Constitutive Rules*, Marcello Ceci, Tom Butler, Leona

O'Brien and Firas Al Khalil illustrate a heuristic approach for the representation of alethic statements as part of a methodology aimed at ensuring effective translation of the regulatory text into a machine-readable language. The methodology includes an intermediate language, accompanied by an XML persistence model, and introduces a set of "legal concept patterns" to specifically represent the different constitutive statements that can be found in e.g. financial regulations. In *An Architecture for Establishing Legal Semantic Workflows in the Context of Integrated Law Enforcement*, Markus Stumpner, Wolfgang Mayer, Pompeu Casanovas and Louis de Koker develop a federated data platform that aims to enable the execution of integrated analytics on data accessed from different external and internal sources, and to enable effective support of an investigator or analyst working to evaluate evidence and manage investigation structure. By preventing the shortcomings of traditional approaches, e.g. high costs and silos-effects, the chapter also aims to show how this integration can be compliant. In *Contributions to Modelling Patent Claims when Representing Patent Knowledge*, Simone Reis, Andre Reis, Jordi Carrabina and Pompeu Casanovas examine the modelling of patent claims in ontology based representation of patent information. They relate to the internal structure of the claims and the use of the all- element rule for patent coverage, in order to offer the general template for the structure of the claim, and provide the visualization of the claims, the storage of claim information in a web semantics framework, and the evaluation of claim coverage using Description Logic. In *Execution and Analysis of Formalized Legal Norms in Model Based Decision Structures*, Bernhard Walzl, Thomas Reschenhofer and Florian Matthes describe a decision support system to represent the semantics of legal norms, whereas a model based expression language (MxL) has been developed to coherently support the formalization of logical and arithmetical operations. Such legal expert system is built upon model based decision structures and three different components, namely a model store, a model execution component, and an interaction component, have been worked out, so as to finally test the execution and analysis of such structured legal norms vis-à-vis the German child benefit regulations. In *Causal Models of Legal Cases*, Ruta Liepina, Giovanni Sartor and Adam Wyner draw the attention to the requirements for establishing and reasoning with causal links. In light of a semi-formal framework for reasoning with causation that uses strict and defeasible rules for modelling factual causation in legal cases, the chapter takes into account the complex relation between formal, common sense, norm and policy based considerations of causation in legal decision making with particular focus on their role in comparing alternative causal explanations. In *Developing Rule-Based Expert System for People with Disabilities*, Michał Araszkiewicz and Maciej Klodawski present the features of a moderately simple legal expert system devoted to solving the most frequent legal problems of disabled persons in Poland. By casting light on the structure of the expert system and its methodology, the succession law of Poland and its procedures delivers sufficient material to reveal the most important issues concerning such a project on a rule-based expert system.

5.3 Legal Vocabularies and Natural Language Processing

The third part of the volume regards legal vocabularies and natural language processing. The eight contributions include EuroVoc-based Summarization of European Case Law, in which Florian Schmedding illustrates the on-going development of a multilingual pipeline for the summarization of European case law. By applying the TextRank algorithm on concepts of the EuroVoc thesaurus, so as to extract summarizing keywords and sentences, the intent is to demonstrate the feasibility and usefulness of the presented approach for five different languages and 18 document sources. In *Aligning Legivoc Legal Vocabularies by Crowdsourcing*, Hughes-Jehan Vibert, Benoit Pin and Pierre Jouvelot present the first Internet-based platform dedicated to the diffusion, edition and alignment of legal vocabularies across countries. As a seamless path for governments to disseminate their legal foundations and specify semantic bridges between them, the chapter describes the general principles behind the legivoc framework while providing some ideas about its implementation, e.g. crowdsourcing the alignment of legal corpora together. In *Data Protection in Elderly Health Care Platforms*, Ângelo Costa, Aliaksandra Yelshyna, Teresa C. Moreira, Francisco Andrade, Vicente Julián and Paulo Novais deal with solutions to the increasing cognitive problems that affect the elderly population in light of the iGenda project, which aims to build safe environments that adapt themselves to one's individual needs through a Cognitive Assistant inserted in the Ambient Assisted Living area. Whereas one of the main issues concerns the protection of the data flowing within the system and the protection of the user's fundamental rights, the chapter clarifies the principles and legal guarantees of data protection, embracing appropriate solutions for technological features that may be a threat. In *Assigning Creative Commons Licenses to Research Metadata: Issues and Cases*, Marta Poblet, Amir Aryani, Paolo Manghi, Kathryn Unsworth, Jingbo Wang, Brigitte Hausstein, Sunje Dallmeier-Tiessen, Claus-Peter Klas, Pompeu Casanovas and Víctor Rodríguez-Doncel tackle the problem of lack of clear licensing and transparency of usage terms and conditions for research metadata. Making research data connected, discoverable and reusable are the key enablers of the new data revolution in research. Accordingly, the chapter does not only discuss how the lack of transparency can hinder discovery of research data and make it disconnected from publication and other trusted research outcomes. In addition, the chapter suggests the application of Creative Commons licenses for research metadata, and provides some examples of the applicability of this approach to internationally known data infrastructures. In *Dataset Alignment and Lexicalization to Support Multilingual Analysis of Legal Documents*, Armando Stellato, Manuel Fiorelli, Andrea Turbati, Tiziano Lorenzetti, Peter Schmitz, Enrico Francesconi, Najeh Hajlaoui and Brahim Batouche tackle the complexity of the EU legal system, in which both the linguistic and the conceptual aspects mutually interweave into a knowledge barrier that is hard to break. In order to create a platform for multilingual cross-jurisdiction accessibility to legal content in the EU, the chapter addresses the challenge of Semantic Interoperability at both the conceptual and lexical level, by developing a coordinated set of instruments for advanced lexicalization of RDF resources (be them ontologies, thesauri and datasets in general) and for alignment of their content. In *A Multilingual Access Module to Legal Texts*, Kiril Simov, Petya Osenova, Iliana Simova, Hristo

Konstantinov and Tenyo Tyankov introduce a Multilingual Access Module, which translates the user's legislation query from its source language into the target language, and retrieves the detected texts that match the query. More particularly, the unit consists of two sub modules, i.e. an Ontology-based and a Statistical Machine Translation units, which have their own drawbacks, so that both are used in an integrated architecture, in order to profit from each other. In Combining Natural Language Processing Approaches for Rule Extraction from Legal Documents, Mauro Dragoni, Serena Villata, Williams Rizzi and Guido Governatori address the problem of moving from a natural language legal text to the respective set of machine-readable conditions, by combining the linguistic information provided by WordNet and a syntax-based extraction of rules from legal texts, with a logic-based extraction of dependencies between chunks of such texts. Such a combined approach leads to a powerful solution towards the extraction of machine-readable rules from legal documents, which is evaluated over the Australian "Telecommunications consumer protections code". In Analysis of Legal References in an Emergency Legislative Setting, Monica Palmirani, Luca Cervone, Ilaria Bianchi and Francesco Draicchio provide for a taxonomy of legal citations that set up an interesting apparatus for analysing a country's legislative approach. By investigating the references of a legal corpus of ordinances issued by the Regional Commissioner for Emergency and Reconstruction over the first eighteen months after the 2012 earthquake in Emilia-Romagna, the chapter scrutinizes the critical issues arising in the regulative strategy for emergency situations. By distinguishing groupings based on lexical-textual analysis and groupings based on structural element, the aim is to help lawmakers act better in future disasters, to extract information concerning the number and the types of modifications produced, and to support the debate on emergency national laws that deal with natural disasters.

5.4 Legal Ontologies and Semantic Annotation

As to the fourth part of this volume on legal ontologies and semantic annotation, it comprises eight chapters. In Using Legal Ontologies with Rules for Legal Textual Entailment, Biralatei James Fawei, Adam Wyner, Martin Kollingbaum and Jeff Z. Pan describe an initial attempt to model and implement the automatic application of legal knowledge using a rule-based approach. Whilst an NLP tool extracts information to instantiate an ontology relative to concepts and relations, ontological elements are associated with legal rules written in SWRL to draw inferences to an exam question. Although further development of the methodology and identification of key issues require future analysis, the preliminary results on a small sample are promising. In KR4IPLaw Judgment Miner- Case-Law Mining for Legal Norm Annotation, Shashishekar Ramakrishna, Łukasz Górski and Adrian Paschke offer a proof-of-concept implementation for automatizing the process of identifying the most relevant judgments pertaining to a legal field and further transforming them into a formal representation format. On this basis, the annotated legal section and its related judgments can be mapped into a decision model for further down the line processing. In Conceptual Annotation of Legal Documents with Ontology Concepts, Kolawole John Adebayo, Luigi Di Caro and Guido Boella illustrate a novel task of semantic labelling, which exploits ontology in providing a fine-grained conceptual document segmentation and

annotation. By dividing documents into semantically-coherent blocks and performing conceptual tagging for efficient information filtering, the chapter presents a promising solution, since the proposed task has several applications, such as granular information filtering of legal texts, text summarization and information extraction among others. In *Reuse and Reengineering of Non-ontological Resources in the Legal Domain*, Cristiana Santos, Pompeu Casanovas, Víctor Rodríguez Doncel and Leon van der Torre explain the processes of reusing and reengineering available non-ontological resources in the legal domain into ontologies, taking into account their specificities, as they are highly heterogeneous in their data model and contents. The description of these processes is further clarified by using a case-study in the consumer law domain. In *Ontology Modelling for Criminal Law*, Chiseung Soh, Seungtak Lim, Kihyun Hong and Young- Yik Rhim propose a general methodology for designing criminal law ontologies and rules. The chapter illustrates the super-domain ontology that contains the common characteristics of criminal law, in order to explain the rule design method of criminal law and present the application of the anti-graft act in Korea as an example. In *Con- trattiPubblici.org*, a Semantic Knowledge Graph on Public Procurement Information, Giuseppe Futia, Federico Morando, Alessio Melandri, Lorenzo Canova and Francesco Ruggiero present the *ContrattiPubblici.org* project, which aims to develop a semantic knowledge graph based on linked data principles in order to overcome the fragmentation of existent datasets, to allow easy analysis, and to enable the reuse of information. The objectives are to increase public awareness about public spending, to improve transparency on the public procurement chain, and to help companies retrieving useful knowledge for their business activities. In *Application of Ontology Modularization for Building a Criminal Domain Ontology*, Mirna El Ghosh and Habib Abdulrab carry out a survey on ontology modularization and present a modular approach to build a criminal modular domain ontology (CriMonto) for modelling the legal norms of the Lebanese criminal system. CriMonto, which will also be used for a legal reasoning system, is composed of four independent modules that will be combined together to compose the whole ontology. In *A Linked Data Terminology for Copyright Based on Ontolex-Lemon*, Víctor Rodríguez-Doncel, Cristiana Santos, Pompeu Casanovas, Asunción Gómez Pérez and Jorge Gracia discuss Ontolex-lemon, namely, the de facto standard to represent lexica relative to ontologies that can be used to encode term banks as RDF. A multi-lingual, multi-jurisdictional term bank of copyright-related concepts has been published as linked data based on the ontolex-lemon model. The terminology links information from several sources and the terms have been hierarchically arranged, spanning multiple languages and targeting different jurisdictions. Whilst the term bank has been published as a TBX dump file and is publicly accessible as linked data, it has been used to annotate common licenses in the RDFLicense dataset.

5.5 Legal Argumentation

Five chapters compose the fifth part of this volume on legal argumentation. In *Abstract Agent Argumentation (Triple-A)*, Ryuta Arisaka, Ken Satoh and Leon van der Torre introduce a Dung style theory of abstract argumentation, which they call triple-A, in which each agent decides autonomously whether to accept or reject her own arguments.

By distinguishing between trusted arguments, selfish agents, and social agents, the extensions of globally accepted arguments are defined using a game theoretic equilibrium definition. In *A Machine Learning Approach to Argument Mining in Legal Documents*, Prakash Poudyal analyzes and evaluates the natural language arguments present in the European Court of Human Right (ECHR) Corpus. By dividing the research into four modules, work on argumentative sentences vs. non-argumentative sentences in narrative legal texts, is accomplished, so as to flesh out the features of this module and conduct an experiment in Sequential Optimization Algorithm and Random Forest Algorithm, which can be used as the basis of a general argument mining framework. In *Answering Complex Queries on Legal Networks: a Direct and a Structured IR Approaches*, Nada Mimouni, Adeline Nazarenko and Sylvie Salotti compare two methods of search in legal collection networks, so as to present new functionalities of search and browsing. Relying on a structured representation of the collection graph, the first approach allows for approximate answers and knowledge discovery, whilst the second one supports richer semantics and scalability but offers fewer search functionalities. As a result, the chapter indicates how those approaches could be combined to get the best of both. In *Inducing Predictive Models for Decision Support in Administrative Adjudication*, Karl Branting, Alexander Yeh and Brandy Weiss explore the hypothesis that predictive models induced from previous administrative decisions can improve subsequent decision-making processes. In light of three different datasets, three different approaches for prediction in their domains were tested, showing that each approach was capable of predicting outcomes. By exploring several approaches that use predictive models to identify salient phrases in the predictive texts, the chapter proposes a design for incorporating this information into a decision-support tool. In *Arguments on the Correct Interpretation of Sources of Law*, Robert van Doesburg and Tom van Engers deal with the formalization of legal reasoning and the representation of law through computational models of argumentation. Whereas most examples presented in literature can be characterized as post-hoc theory construction, the chapter aims to provide an instrument that can be used to inform legal experts on relevant issues in the process of solving current cases, i.e. using the interpretations of legal sources ex-ante. An actual case that is in discussion in the Dutch Tax Administration, in court as well as in Parliament, helps to further clarify this approach.

5.6 Courts, Adjudication and Dispute Resolution

The sixth part of the volume is devoted to courts, adjudication and dispute resolution, which are the subject matter of five chapters. In *Dynamics of the Judicial Process by Defeater Activation*, Martin Moguillansky and Guillermo Simari illustrate a novel activating approach to Argument Theory Change (ATC) for the study of the dynamics of the judicial process. By considering the sentences of two different real criminal procedures, the aim is to contribute to the discussion of how to deal with circumstances of the judicial process like hypothetical reasoning for conducting investigations of a legal case, and for handling the dynamics of the judicial process. In *Claim Detection in Judgments of the EU Court of Justice*, Marco Lippi, Francesca Lagioia, Giuseppe Contissa, Giovanni Sartor and Paolo Torroni address recent approaches to argumentation mining in juridical documents, so as to present two distinct contributions. The

first one is a novel annotated corpus for argumentation mining in the legal domain, together with a set of annotation guidelines. The second one is the empirical evaluation of a recent machine learning method for claim detection in judgments. Whereas the latter method has been applied to context-independent claim detection in other genres such as Wikipedia articles and essays, the chapter shows that this method also provides a useful instrument in the legal domain, especially when used in combination with domain-specific information. In *A Non-intrusive Approach to Measuring Trust in Opponents in a Negotiation Scenario*, Marco Gomes, John Zeleznikow and Paulo Novais propose a threefold approach to trust, that regards the possibility of measuring trust based on quantifiable behaviour, the use of Ambient Intelligence techniques that use a trust data model to collect and evaluate relevant information based on the assumption that observable trust between two entities (parties) results in certain typical behaviours and, finally, relational aspects of trust and parties' conflict styles based on cooperativeness and assertiveness. The main contribution of this chapter is the identification of situations in which trust relationships influence the negotiation performance. In *Network, Visualization, Analytics. A Tool Allowing Legal Scholars to Experimentally Investigate EU Case Law*, Nicola Lettieri, Sebastiano Faro, Delfina Malandrino, Margherita Vestoso and Armando Faggiano dwell on the intersection between Network Analysis (NA), visualization techniques and legal science research questions. Their aim is to bring the network approach into "genuinely legal" research questions, and to create tools that allow legal scholars with no technical skills to make experiments with NA and push new ideas both in legal and NA science, so as to use NA and visualization in their daily activities. In *Electronic Evidence Semantic Structure: Exchanging Evidence across Europe in a Coherent and Consistent Way*, Maria Angela Biasiotti and Fabrizio Turchi provide for a seminal work on a common and shared understanding of what Electronic Evidence is and how it should be treated in the EU context and in the EU member states. The chapter develops a tailor-made categorization of relevant concepts which provides a starting analysis for the exchange of Electronic Evidence and data between judicial actors and LEAs, with a specific focus on issues of the criminal field and criminal procedures. This semantic structure might represent a good starting point for the alignment of electronic evidence concepts all over Europe in a cross border dimension.

Acknowledgments. Law and Policy Program of the Australian government funded Data to Decisions Cooperative Research Centre (<http://www.d2dcre.com.au/>); Meta-Rule of Law DER2016-78108-P, Research of Excellence, Spain. This work was also partially supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 690974 "MIREL: MIning and REasoning with Legal texts".

References

1. Regalado, A.: Google's AI explosion in one chart. MIT technology review, 25 March 2017. <https://www.technologyreview.com/s/603984/googles-ai-explosion-in-one-chart/>
2. Pagallo, U., Massimo Durante, M., Monteleone, S.: What is new with the internet of things in privacy and data protection? Four legal challenges on sharing and control in IoT. In: Leenes, R., van Brakel, R., Gutwirth, S., de Hert, P. (eds.) Data Protection and Privacy: (In) visibilities and Infrastructures. LGTS, pp. 59–78. Springer, Dordrecht (2017)
3. Susskind, R., Susskind, D.: The Future of the Professions. Oxford University Press, Oxford (2015)
4. Abelson, H., et al.: Keys under the doormat: mandating insecurity by requiring government access to all data and communications. MIT Computer Science and AI Laboratory Technical report, 6 July 2015
5. Abelson, H., et al.: The future of life institute, an open letter: research priorities for robust and beneficial artificial intelligence (2015). <http://futureoflife.org/ai-open-letter/>. Accessed 18 Oct 2016
6. IEEE Standards Association: The Global Initiative for Ethical Considerations in the Design of Autonomous Systems, Forthcoming (2017)
7. Hart, H.L.A.: The Concept of Law. Clarendon, Oxford (1961)
8. Leenes, R., Lucivero, F.: Laws on robots, laws by robots, laws in robots: regulating robot behaviour by design. Law Innov. Technol. 6(2), 193–220 (2014)
9. Garfinkel, S., Spafford, G.: Web Security and Commerce. O'Reilly, Sebastopol (1997)
10. Pagallo, U.: The legal challenges of big data: putting secondary rules first in the field of EU data protection. Eur. Data Protect. Law Rev. 3(1), 34–46 (2017)
11. Pagallo, U.: LegalAIze: tackling the normative challenges of artificial intelligence and robotics through the secondary rules of law. In: Corrales, M., Fenwick, M., Forgó, N. (eds.) New Technology, Big Data and the Law. Perspectives in Law, Business and Innovation. PLBI, pp. 281–300. Springer, Singapore (2017)
12. Weng, Y.H., Sugahara, Y., Hashimoto, K., Takanishi, A.: Intersection of “Tokku” special zone, robots, and the law: a case study on legal impacts to humanoid robots. Int. J. Soc. Robot. 7(5), 841–857 (2015)
13. Pagallo, U., Durante, M.: The pros and cons of legal automation and its governance. Eur. J. Risk Regul. 7(2), 323–334 (2016)
14. Brody, P., Pureswaran, V.: Device democracy: saving the future of the internet of things. IBM, September 2014
15. Boella, G., Van Der Torre, L., Verhagen, H.: Ten challenges for normative multiagent systems. In: Dagstuhl Seminar Proceedings. Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2008)
16. Hansen, J., Pigozzi, G., Van Der Torre, L.: Ten philosophical problems in deontic logic. In: Dagstuhl Seminar Proceedings. Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2007)
17. Casanovas, P., Palmirani, M., Peroni, S., van Engers, T., Vitali, F.: Semantic web for the legal domain: the next step. Semant. Web 7(3), 213–227 (2016)
18. Andresen, G.: Nexxus Whitepaper (2017). <http://nexxusuniversity.com>
19. Jentzsch, C.: Decentralized autonomous organization to automate governance (2016). <https://download.slock.it/public/DAO/WhitePaper.pdf>. Accessed 23 June 2016
20. Atzei, N., Bartoletti, M., Cimoli, T.: A survey of attacks on ethereum smart contracts (SoK). In: Maffei, M., Ryan, M. (eds.) POST 2017. LNCS, vol. 10204, pp. 164–186. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54455-6_8
21. Xu, J.J.: Are blockchains immune to all malicious attacks? Financ. Innov. 2(1), 25 (2016)

22. Arruñada, B.: Property as sequential exchange: the forgotten limits of private contract. *J. Inst. Econ.* 13(4), 753–783 (2017)
23. DuPont, Q., Maurer, B.: Ledgers and Law in the Blockchain. *Kings Review*, 23 June 2015. <http://kingsreview.co.uk/magazine/blog/2015/06/23/ledgers-and-law-in-the-blockchain>
24. Davidson, S., De Filippi, P., Potts, J.: Economics of blockchain. In: *Proceedings of Public Choice Conference Public Choice, Conference*, May 2016, Fort Lauderdale, United States (2016). <https://doi.org/10.2139/ssrn.2744751>. HAL Id. hal-01382002
25. English, M., Auer, S., Domingue, J.: Block chain technologies and the semantic web: a framework for symbiotic development. In: *Lehmann, J., Thakkar, H., Halilaj, L., Asmat, R. (eds.) Computer Science Conference for University of Bonn Students*, pp. 47–61 (2016)
26. Rodríguez-Doncel, V., Santos, C., Casanovas, P., Gómez-Pérez, A.: Legal aspects of linked data—the European framework. *Comput. Law Secur. Rev.* 32(6), 799–813 (2016)
27. Casanovas, P., Mendelson, D., Poblet, M.: A linked democracy approach for regulating public health data. *Health Technol.* 7(4), 519–537 (2017)
28. Poblet, M., Casanovas, P., Plaza, E. (ed.) *Linked Democracy: Artificial Intelligence for Democratic Innovation*. <http://ceur-ws.org/Vol-1897/>
29. Poblet, M., Casanovas, P., Rodríguez-Doncel, V.: *Linked Democracy*. Springer Briefs, Cham: Springer Nature (2018, forthcoming)
30. Sartor, G., Palmirani, M., Francesconi, E., Biasiotti, M.A.: *Legislative XML for the Semantic Web: Principles, Models, standards for Document Management*. Springer, New York (2011)
31. Giovannini, M.P., Palmirani, M., Francesconi, E.: Linee guida per la marcatura dei documenti normativi secondo gli standard Normainrete, p. 200. European Press Academic Publishing, Firenze (2012)
32. Boer, A., de Maat, E., Francesconi, E., Lupo, C., Palmirani M., Winkels R.: General XML Format(s) for Legal Sources, Deliverable 3.1, Estrella Project - European project for Standardized Transparent Representations in order to Extend Legal Accessibility, Proposal/Contract no.: 027655, European Commission (2007)
33. Breuker, J.A.P.J., et al.: OWL ontology of basic legal concepts (LKIF-Core). Estrella: Deliverable 1.4., AMSTERDAM, UVA 2007, pp. 138 (2007)
34. Gordon, Thomas F.: Constructing legal arguments with rules in the legal knowledge interchange format (LKIF). In: *Casanovas, P., Sartor, G., Casellas, N., Rubino, R. (eds.) Computable Models of the Law. LNCS (LNAI)*, vol. 4884, pp. 162–184. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85569-9_11
35. Vitali F., Palmirani M.: Akoma Ntoso Release Notes. Accessed 5 April 2018
36. Opijnen, M., Palmirani, M., Vitali, F., Agnoloni, T.: Towards ECLI 2.0. In: *2017 International Conference for E-Democracy and Open Government*, Los Alamitos, CA, IEEE, 2017, P6082, pp. 1–9 (2017)
37. ELI Task Force: Technical ELI implementation guide citations, Publications Office of the European Union, Luxembourg (2015). <https://doi.org/10.2830/74251>
38. Athan, T., Governatori, G., Palmirani, M., Paschke, A., Wyner, A.: LegalRuleML: design principles and foundations. In: *Faber, W., Paschke, A. (eds.) Reasoning Web 2015. LNCS*, vol. 9203, pp. 151–188. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21768-0_6
39. Peroni, S., Palmirani, M., Vitali, F.: UNDO: the united nations system document ontology. In: *d’Amato, C., Fernandez, M., Tamma, V., Lecue, F., Cudré-Mauroux, P., Sequeda, J., Lange, C., Heflin, J. (eds.) ISWC 2017. LNCS*, vol. 10588, pp. 175–183. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68204-4_18

40. Barabucci, G., Cervone, L., Palmirani, M., Peroni, S., Vitali, F.: Multi-layer markup and ontological structures in Akoma Ntoso. In: Casanovas, P., Pagallo, U., Sartor, G., Ajani, G. (eds.) AICOL - 2009. LNCS (LNAI), vol. 6237, pp. 133–149. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16524-5_9
41. Casalicchio, E., Cardellini, V., Interino, G., Palmirani, M.: Research challenges in legal-rule and QoS-aware cloud service brokerage. *Future Gen. Comput. Syst.* 78, 211–223 (2016)
42. Governatori, G., Hashmi, M., Lam, H.-P., Villata, S., Palmirani, M.: Semantic business process regulatory compliance checking using LegalRuleML. In: Blomqvist, E., Ciancarini, P., Poggi, F., Vitali, F. (eds.) EKAW 2016. LNCS (LNAI), vol. 10024, pp. 746–761. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-49004-5_48
43. Ashley, K.D.: *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*, p. 446. Cambridge University Press, Cambridge (2017)
44. Palmirani, M., Vitali, F.: Legislative drafting systems. In: *Usability in Government Systems*. Morgan Kaufmann, New York, pp. 133–151 (2012)
45. Rossi, A., Palmirani, M.: A visualization approach for adaptive consent in the European data protection framework. In: *2017 International Conference for E-Democracy and Open Government*, Los Alamitos, CA, IEEE, 2017, pp. 1–12 (2017)
46. Haapio, H., Hagan, M., Palmirani, M., Rossi, A.: Legal design patterns for privacy. In: *Data Protection/LegalTech Proceedings of the 21st International Legal Informatics Symposium IRIS 2018*, II, pp. 445–450. Editions Weblaw, Bern (2018)