IRIS AperTO

UNIVERSITÀ DEGLI STUDI DI TORINO

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

The pros and cons of legal automation and its governance

This is the author's manuscript

Original Citation:

Availability:

This version is available http://hdl.handle.net/2318/1640462 since 2017-05-31T12:11:11Z

Published version:

DOI:10.1017/S1867299X00005742

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

25 April 2024

The Pros and Cons of Legal Automation and its Governance

Ugo Pagallo and Massimo Durante

Professors of Jurisprudence, Law School, University of Torino,
Lungo Dora Siena 100 A, 10153 Torino, Italy
ugo.pagallo@unito.it, massimo.durante@unito.it

Abstract. The paper examines the sector of legal automation, its advantages and drawbacks, the ways in which legal constraints and safeguards can be embedded into technology and how the law may govern such an aim to regulate human behaviour through codes, IT architectures, and design. By stressing both advantages and drawbacks of legal automation, the paper does not intend to suggest that the latter is something "neutral." Rather, legal automation should be conceived as a set of constraints and affordances that transform, or reshape, the environment of people's interaction and more specifically, the interplay of human and artificial agents, thereby affecting basic pillars of the (rule of) law. The overall aim is to flesh out goals and values that are at stake with choices of technological dependence, delegation and trust, so as to determine the good mix between legal automation and public deliberation.

Keywords: Automation, Delegation, Governance, Legal Ontologies, Normative Context, Privacy by Design, Trust, Self-enforcing Technologies.

1 Introduction

This paper examines the sector of legal automation, its advantages and drawbacks, the ways in which legal constraints and safeguards can be embedded into technology and how the law may govern such an aim to regulate human behaviour through codes, IT architectures, and design. On the one hand, the pros of legal automation have been stressed time and again by experts of such fields as AI and the law, legal informatics, and more. Here, the intent is to improve the efficiency, consistency, comprehensibility, and predictability of legal and judicial systems through, for example, machine learning and data mining techniques for legal applications, computational models of legal reasoning, ICT applications to support the legal domain, and more. On the other hand, the cons of legal automation are the bread and butter of work on the new surveillance society, the death of privacy, legal regulation by design, codes, and so forth, that shed light on the challenges of self-enforcing technologies, the implicit values of technology and its invisibility, down to the lack of public debate brought on by automation, as occurs with the use of drones on the battlefield with no parliamentary authorization. The emphasis on both advantages and drawbacks of legal automation, however, does not intend to suggest that the latter is something "neutral," namely a simple means to achieve whatsoever end. Rather, legal automation should be conceived as a set of constraints and affordances (Durante 2011a) that transform, or reshape, the environment of people's interaction and more specifically, the interplay of human and artificial agents, thereby affecting basic pillars of the (rule of) law. The information revolution and the new scenario of societies that depend on information as a vital resource, have already impacted crucial aspects of legal representation and resolution: think about the right of the individuals to have a say in the decisions affecting them on the internet.

Some, both public and private actors, have proposed to tackle the complexity of this new situation with the same tools of automation. Over the past 20 years, we have had several instances of this sort of meta-automation: for example, contemplate Article 8 of WIPO's 1996 *Copyright Treaty* and Article 14 of the twin *Performances and Phonograms Treaty*, i.e. the legal umbrella for the adoption of such automatic techniques as digital right management (DRM), which enable copyright holders to monitor and regulate the use of their protected artefacts. Likewise, reflect on the use of automatic systems for filtering all electronic communications, such as those proposed by the EU Commission in the 2010 Report on the copyright directive (SEC-2010-1589 final), and adopted in UK with the Digital Millennium Act from the same year. Whereas the idea is also popular in the field of privacy by design, several projects for the development of distributed computer systems capable of self-management have been developed, e.g. the IBM's autonomic computing project started in 2001 (Durante 2011b). By insisting on such trends, the paper does not intend to take

sides between the Scylla and Charybdis of legal automation and public deliberation as a zero sum game. On the contrary, the aim is to flesh out 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.

At times, we have to admit that legal automation is not only unproblematic but positive, since it helps to improve the functioning of legal and judicial systems in terms of foreseeability and legal certainty, efficiency, etc. Yet, we should be prepared for the other way around. Legal automation may not only infringe rules and principles of the law, as the EU Court of Justice's decision in *Netlog* (C-360/2010) illustrates. In addition, legal automation has already raised several hard cases in today's laws of war (e.g. the use of autonomous lethal weapons on the battlefield), business law (e.g. matters of liability for contracts executed by robots and other artificial agents), and tort law (e.g. forms of distributed responsibility that hinge on multiple accumulated actions of both humans and computers that make difficult to ascertain what is, or should be, the information content of the corporate entity so as to determine its accountability). As a result, how should we tackle these legal cases? Is there a unique right answer out there or conversely, should we find a reasonable balance, or even compromise, between many conflicting interests (Durante 2013) that cannot be subject to legal automation?

In order to offer a hopefully comprehensive view on these issues, the paper is divided into three parts. Section 2 explores the pros of legal automation and more specifically, a case study that concerns a sub-field of AI & the law, such as legal ontologies (Section 2.1). How the latter is implemented in the legal field is then elucidated with the example of the principle of privacy by design (Section 2.2). Some automatic versions of this principle are furthered in Section 2.3, so as to introduce the second part of this paper. Section 3 has in fact to do with the cons of legal automation and how the latter may impact both the requirements and functions of the law, i.e. what the law is supposed to be (requirements), and what it is called to do (functions). In particular, focus is on the development of self-enforcing technologies, such as the use of automatic filtering systems on the internet (Section 3.1). Although such techniques are quite problematic and even illegal (Section 3.2), it is noteworthy how popular they are among national sovereigns, governance actors, and private companies alike (Section 3.3). The final part of the paper draws the attention to the formation and stewardship of the formal and informal rules that regulate this crucial sector of current legal systems (Section 4). The governance of legal automation is deepened through the analysis of both the internal and external limits of legal automation (Section 4.1), along with the theoretical framework of the interaction between law and technology as competing regulatory systems (Section 4.2), the need for an institutional forum for deliberation (Section 4.3), with ends up with the distinction between plain and hard cases of the law (Section 4.4). At the end of the day, since the legal hard cases of automation will probably be the main subject of lawyers, scholars, and policy makers for quite a long time, the aim of this paper is to provide a normative stance with which to tackle such legal hard cases.

2 The Pros of Legal Automation

The grandfather of today's AI & the law is the German philosopher Gottfried Wilhelm Leibniz. As shown by his early work on jurisprudence, such as his *magister* paper on *Specimen quaestionum philosophicarum ex jure collectarum* (1664), his *Habilitation* research known as *Disputatio de arithmetica de complexionibus* (1665), followed by *De casibus perplexis in iure* (1666), and moreover *Dissertio de arte combinatoria* (1666), down to the summa *Nova methodus* from 1667, Leibniz's overall ambition was to turn legal arguments into computing through combinatorial analysis, probability calculus, and binary arithmetic (Pagallo 2006). In the words of *de Arte Combinatoria*, "there would be no more need of disputation between two philosophers than between two accountants. For it would suffice to take their pencils in their hands, and say to each other: Let us calculate."

Three and a half centuries later, Leibniz's vision has partially come true. A simple check of the websites and programs of such international conferences as *Artificial Intelligence and Law* (ICAIL), *Legal Knowledge and Information System* (Jurix), *AI and Legal Complexity* (AICOL), etc., would confirm this point (Casanovas et al. 2014 and 2010; Palmirani et al. 2012). In order to illustrate how the automation of legal reasoning can improve the efficiency, comprehensibility, consistency,

and predictability of legal and judicial systems, let us restrict the focus of the analysis on a sub-set of today's research of AI & the law, such as legal ontologies. Once familiar with this field (Section 2.1), and how it works in the legal domain of data protection (Section 2.2), limits and risks of this implementation should be clear (Section 2.3). At that point we'll be ready for the other side of the coin: the cons of legal automation.

2.1 Legal Ontologies

Legal ontologies is the field of artificial intelligence (AI) that aims to model concepts traditionally employed by lawyers through the formalization of norms, rights, and duties, in fields like criminal law, administrative law, civil law, etc. (Breuker et al. 2009; Casanovas et al. 2010). The objective is that even a machine should comprehend and process this very information, by preliminarily distinguishing between the part of the ontology containing all the relevant concepts of the problem domain through the use of taxonomies (e.g. ontological requirements), and the ontology which includes both the set of rules and restraints that belong to that problem domain (e.g. ontological constraints*)*. An expert system should thus process the information in compliance with regulatory legal frameworks through the conceptualization of classes, relations, properties, and instances pertaining to a given problem domain. In technical terms, we should pay attention to the bottom-up approach that starts from legal concepts defined by scholars. A traditional top-down approach works well for the topmost level, where the representation instruments are at the disposal of the ontology-builders and the basic conceptual primitives such as relation, role, qualia, processes, etc., are precisely defined. However, a lot of issues arise when the core ontology level is taken into account, because the amount of information involved in the project of making legal information automatic is hardly compressible. Simply put, many regulations not only include "top normative concepts" such as notions of validity, obligation, prohibition, and the like. These rules present also highly context-dependent normative concepts like, for example, in the field of data protection, notions of personal data, security measures, or data controllers (Pagallo 2011a).

This difficulty does not mean, of course, that work on legal ontologies should be abandoned. On the contrary, these problems suggest a bottom-up rather than a top-down approach, in order to lawfully process growing amounts of data. By splitting the work into several tasks and assigning each to a working team, we should start from smaller parts and sub-solutions of the project, to end up with global answers. The evaluation phase consists in testing the internal consistency of the project and, according to that which Herbert Simon used to dub as the "generator test-cycle," the evaluation entails the decomposition of the complete design into functional components. The test generates alternatives and examines them against the set of requirements and constraints, so that "important indirect consequences will be noticed and weighed. Alternative decompositions correspond to different ways of dividing the responsibilities for the final design between generators and tests" (Simon 1996: 128). Further criteria and empirical methods have been proposed: apart from functional efficiency, consider the robustness, reliability, and usability of projects on legal automation. Their evaluation and verification can additionally employ automated and regression-oriented tests, use of prototypes, internal checks among the design team, users tests in controlled environments, surveys, interviews, and more (Pagallo 2012).

On this basis, we can quantify the growing amount of data processed in compliance with regulatory frameworks, as occurs with several projects for representing, processing and retrieving legal information in, say, large databases, through analogical legal arguments, or via document modelling (Sartor et al. 2011). Likewise, consider work on legal ontologies for the support of privacy preservation in location-based services (Mitre et al. 2006), the management of information systems (Abou-Tabir and Berlik 2006; Casellas 2011), or middleware architectures for data protection (Lioukadis et al. 2007), each of which aims at integrating smaller parts and sub-solutions into the design of the project. Remarkably, there are even cases where the conceptualization of classes, relations, properties, and instances pertaining to a given problem domain, does not seem particularly complex, e.g. the design of information systems for hospitals to ensure that patient names are kept separated from data on medical treatments or health status (Pagallo 2011a). In general terms, the overall idea is to embed legal constraints, or safeguards, into information systems and other types of technology, so as to automatically abide by the rules and principles of current legal frameworks. Let us now explore how far this approach goes in the field of data protection.

2.2 Two Roads to Automatic Privacy by Design

The idea of embedding privacy safeguards in information systems and other types of technology is nothing new. Although the Ontario's Privacy Commissioner, Ann Cavoukian, invented the formula "privacy by design" in the late 1990s, the notion can be found in recital 46 and article 17 of the EU directive 46 from 1995 (D-46/95/EC), according to which "Member States shall provide that the controller must implement appropriate technical… measures to protect personal data." More recently, the formula appears in articles 23 and 30 of the EU Commission's proposal for a new data protection regulation from January 2012, much as in § 3.4.4.1 of the document with which the Commission illustrated the proposal. In the wording of the EU Parliament's amendment 118 from March 2014, privacy by design refers to "comprehensive procedural safeguards regarding the accuracy, confidentiality, integrity, physical security and deletion of personal data."

Admittedly, the idea of embedding privacy safeguards into technology may not entail any automation. For example, privacy by design may aim to encourage individual change of conduct by broadening the range of choices available through the configuration of user-friendly setting options, or the privacy controls of IT interfaces and default mechanisms (Pagallo 2012). Consider the ubiquitous commons browser plugin ("UCBP") scheme, through which data, information and knowledge can be encrypted before going into the servers of a service provider, depending on the decisions of the data subject on how to modulate different levels of access to, and control over, such data and information (Iaconesi 2014). By using UCBP to share the decryption keys for the data on a peer-to-peer (p2p) network external to the service provider, users may determine the license for the content of such information with some metadata. Hosted in the servers of the service provider, data and information become usable only once subjects or entities wanting to use it has verified they can, on the basis of their identity, context and intended usage scheme. Consequently, the service provider's databases appear as a storage space of securely encrypted data which can only be accessed if the entity desiring to use it, is included in the license distributed with the decryption key, and which is available in the p2p network. According to the EU Art. 29 Working Party's opinion on "the future of privacy" from 1st December 2009, the principle of privacy by design can thus be implemented in a bottom-up way, i.e. grounded on the autonomous choices of individuals through self-regulation and competition among private organizations (WP 168).

Still, the principle of privacy by design can also be understood in a fully automated way. Here, we have to further distinguish a field-dependent approach from an ideological stance. In the first case, the aim to make privacy safeguards automatic hinges on the specific problems with which we are dealing, and that partially overlap with work on legal ontologies mentioned above in the previous section. Reflect on the field of robotics and more particularly, the set of data protection and privacy issues raised by such a sub-class of service robots, as domestic or consumer robots, that either suggest the regulation of user behaviour through the design of the artificial agent, that is, by designing robots in such a way that unlawful actions of humans are not allowed, or the regulation of robot behaviour through design, that is, by embedding normative constraints into the design of the artificial agent (Leenes and Lucivero 2014). Some legal safeguards, such as data security through encryption and data access control, can be embedded into the software and interface of the robot. Likewise, "requirements such as informed consent can be implemented in system design, for example through interaction with users displays and input devices" (RoboLaw 2014). Furthermore, robots could be designed in a privacy-friendly way, so that the amount of data to be collected and processed is reduced to a minimum and in compliance with the finality principle. This means that, pursuant to, say, Article 6(1)(b) of the EU data protection directive 46 from 1995, robots should collect data only insofar as it is necessary to achieve a specified and legitimate purpose.

On the other hand, the ideological approach to the automatic version of privacy by design can be illustrated with Ann Cavoukian's work. Here, regardless of the technology or business practices involved, the overall idea is to view data protection safeguards in proactive rather than reactive terms, that is, making privacy by design preventive and not simply remedial. It follows that personal data should be automatically protected in every information system as its default position, so that, by embedding data protection safeguards into design, a cradle-to-grave, start-to-finish, or

end-to-end lifecycle protection ensures that privacy safeguards are at work even before a single bit of information has been collected (Cavoukian 2010). Whereas the design project should make data protection mechanisms visible and transparent to both IT users and providers, the full functionality of the principle would allow a positive-sum, or win-win game, making trade-offs unnecessary (e.g. privacy vs. security). Moreover, in the words of Cavoukian, the principle "requires architects and operators to keep the interests of the individual uppermost by offering such measures as strong privacy defaults, appropriate notice, and empowering user-friendly options" (*op. cit.*). And yet, notwithstanding such an individual-focused respect for user privacy, is this automatic version of the principle technically feasible, and even desirable?

After all, what may make sense and properly fit the field of roboprivacy by design, can be quite problematic when design applies to human behaviour. Personal choices play indeed a key role when individuals modulate different levels of access and control over their own information, depending on the context and its circumstances. If there is no need to humanize our robotic applications, we should not robotize human life either. This difference is deepened in the next section, which ideally introduces the second part of this paper on the cons of legal automation.

2.3 The Man and the Machine

There is a number of ethical, legal, and technical reasons why making data protection automatic is problematic. As to ethical reasons, consider how specific design choices may result in conflicts between values and, vice versa, conflicts between values may impact on the features of design: we have evidence that "some technical artefacts bear directly and systematically on the realization, or suppression, of particular configurations of social, ethical, and political values" (Flanagan et al. 2008). In the case of data protection, contemplate the different features that privacy by design acquires, once data protection is grasped in terms of property rights or human dignity, of total control or contextual integrity, of restricted access or limited control over information. At the end of the day, should an artefact be designed in accordance with the traditional European opt-in model for users of electronic communication systems or, vice versa, according to the American opt-out approach? Moreover, reflect upon the information system of hospitals which was mentioned above in section 2.1: should we privilege the efficacy and reliability of that information system in keeping patient names separated from data on medical treatments or health status? But, how about users, including doctors, who may find such mechanism too onerous?

As to the legal reasons against this type of design policy, the development and use of legal automation may curtail both collective and individual autonomy severely. Basic tenets of the rule of law would be at risk, once people's behaviour should unilaterally be determined on the basis of technology (Pagallo 2012). First, there is the threat of updating traditional forms of paternalism through the regulatory tools of technology, because the more personal choices are wiped out by legal automation, the bigger the danger of modelling social conduct via design. Second, attention should be drawn to matters of legal enforcement and its exceptions: what is imperilled here is "the public understanding of law with its application eliminating a useful interface between the law's terms and its application" (Zittrain 2007). Third, rearrangements in the system of legal enforcement are intertwined with redistributions of power and the role of the relevant political institutions with their decisions. As Lawrence Lessig used to warn, the threat is that "controls over access to content will not be controls that are ratified by courts; the controls over access to content will be controls that are coded by programmers" (Lessig 2004).

Finally, the technical difficulties of achieving such a total control through design should be mentioned. Doubts are cast by "a rich body of scholarship concerning the theory and practice of 'traditional' rule-based regulation [that] bears witness to the impossibility of designing regulatory standards in the form of legal rules that will hit their target with perfect accuracy" (Yeung 2007). As stress above in section 2.1, there is indeed the technical difficulty of applying to a machine concepts traditionally employed by lawyers, through the formalization of norms, rights, or duties: legal safeguards do present highly context-dependent notions that raise a number of relevant problems when reducing the informational complexity of a legal system where concepts and relations are subject to evolution. In the words of Bert-Jaap Koops and Ronald Leenes, "the idea of encoding legal norms at the start of information processing systems is at odds with the dynamic

and fluid nature of many legal norms, which need a breathing space that is typically not something that can be embedded in software" (Koops and Leenes 2014).

Making data protection automatic is thus problematic because, in more general terms, legal automation 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). Contrary to the traditional view of the law as a means for social control via a set of rules enforced through the threat of physical sanctions: "if A, then B" (Kelsen 1949), legal automation converts the law into a set of effects (B) that would automatically follow technical instructions (A), rather than sanctions (B) that should follow terms and conditions of legal accountability (A), i.e. that which is, rather than that which should be. The time is ripe to deepen the cons of legal automation.

3 The Cons of Legal Automation

Current advancements of technology have obliged legislators, policy makers, and private actors to forge more sophisticated ways to think about legal enforcement. Moreover, over the past years, several programs and national statutes have been developed by sovereign states to tackle the menace of a new generation of cyber-attacks carried out by other sovereign states or non-state actors. The very beginning of such trends can be traced back to the mid 1990s, when private companies and big business in the media and entertainment sectors had to find out a remedy for the apparent inefficacy of state-action in protecting their own rights in digital environments. While lobbying national and international lawmakers in the copyright field, some of the most relevant companies focused on how to enforce their exclusivity rights through the development of such self-enforcement technologies, as DRM. By enabling right-holders to monitor and regulate the use of their own copyright protected works, companies would have prevented unsolvable issues involving the enforceability of national laws and conflicts of law at the international level, under the umbrella of both Article 8 of WIPO's 1996 *Copyright Treaty* and Article 14 of the twin *Performances and Phonograms Treaty*. Together with severe problems of interoperability and, hence, antitrust-related issues, DRM-compliant systems have nevertheless raised further problems of privacy and data protection, which add to how individuals can be damaged by data mining techniques, behavioural analytics, profiling, etc.

These trends on legal automation are confirmed, and even worsened, by the public sector, especially after the 9/11 scenarios of the "war on terror." Unconventional challenges of cyber-attacks and terrorism have increasingly been testing the framework of legal safeguards that, so far, have represented the salient quality of Western democracies, as shown by the 2013 scandal of the U.S. National Security Agency (NSA)'s Prism project and the UK's GCHQ files.[1] Amongst the forms of legal automation, suffice it to mention the use of self-enforcing technologies, such as filtering systems on the internet that provide a complete packet inspection and Information eXtraction ("IX"), which often go hand in hand with the alleged protection of private rights. For example, in the 2010 Report on the "Application of D-2004/48/EC," i.e. the EU copyright directive, the Commission denounced "the sheer volume and financial value of IP rights infringements… offered by the internet" (SEC-2010-1589 final), so that, in order to stop such illegal activities, internet intermediaries should be obliged to install "a system for filtering all electronic communications" and, especially, p2p applications.

Where all this leads to has been stressed time and again by scholars working on the threats and challenges of the new surveillance society (Lyon 1994; Marx 2002: etc.), the death of privacy (Froomkin 2000), and so forth. Further examples of threats brought on by legal automation can, of course, be made. Think of national laws in the field of data retention, e.g. the EU directive 24 from 2006, much as whether, or to what extent, lethal force should ever be permitted to be fully automated (Pagallo 2013). Here, let us dwell on a particular case regarding the use of filtering systems on the internet (section 3.1); so as to grasp why, at least in the EU, such filtering systems should be deemed as illegal (section 3.2); although, even after the ruling of the Court of Justice in

[1] See John Lanchester, *The Snowden files: why the British public should be worried about GCHQ, The Guardian*, 3 October 2013, at http://www.theguardian.com/world/2013/oct/03/edward-snowden-files-john-lanchester (last accessed 29 August 2014).

2012, such a use is fated to remain an open issue (section 3.3). Then, we will be ready to examine the governance of legal automation (section 4).

3.1 Filtering Information on the Internet

A lively debate over what role internet intermediaries, or service providers ("ISPs"), should have, so as to ensure online security and the protection of individual rights, has occurred in Europe over the past years (Pagallo 2011b). The opinions in the debate can be conceived as falling within the ends of a spectrum that concerns public authorities requiring private companies to safeguard online security, e.g. ISPs as sheriffs of the net and, vice versa, private companies lobbying public authorities to enforce their own rights and interests via the use of filtering systems on the internet. At one end of the spectrum, security trumps civil rights through the use of such filtering systems, because the latter would make impossible any balance between the aim to guarantee online security and the protection of some basic rights, such as data protection, freedom of speech and of information, freedom to conduct a business, and so forth. At the other end of the spectrum, there are constitutional limits to the use of such filtering systems in order to protect some of the basic rights mentioned above. A case discussed before the EU Court of Justice, namely *Netlog* (C-360/10), appears instructive to illustrate the ends of this spectrum.

The plaintiff in *Netlog* was a management company, SABAM, which represents authors, composers, and publishers of musical works in Belgium. As such, SABAM is responsible for authorizing the use by third parties of copyright-protected works of these authors, composers, and publishers. Claiming that a social network, Netlog, made such works available to the public without SABAM's consent and without paying it any fee, the plaintiff thus requested the Court of First Instance in Brussels an injunction against the defendant in order to take appropriate measures to stop the infringement of the plaintiff's intellectual property rights and moreover, to prevent any further infringement. As a result, the national court would have had to issue an injunction against the social network requiring the latter to install a system that, in the wording of the EU Court of Justice, should filter:

a. "Information which is stored on its servers by its service users;
b. Which applies indiscriminately to all of those users;
c. As a preventive measure;
d. Exclusively at its expense; and
e. For an unlimited period;

Which is capable of identifying electronic files containing musical, cinematographic or audio-visual work in respect of which the applicant for the injunction claims to hold intellectual property rights" (C-360/10).

In accordance with the mechanism of the preliminary ruling, the Court of First Instance in Brussels lodged a reference before the EU Court of Justice in Luxembourg, in order to determine rights and duties for processing of information stored on online social networking platforms, and whether introducing a system for filtering that information and prevent files being made available which infringe copyright is lawful in the EU. In addition, the Belgian court asked whether there is a general obligation to monitor stored information. On 16 February 2012, the EU Justices delivered their verdict on whether the use of self-enforcing technologies, such as the filtering system discussed in *Netlog*, is precluded by the EU law.

3.2 Matters of Balance

There are two reasons why the Court of Luxembourg ruled that the filtering system discussed in *Netlog* is precluded by the EU directives on data protection (1995/46/EC), e-commerce (2000/31/EC), copyright (2001/29/EC), and IP (2004/48/EC), much as the freedom to receive or impart information, according to Articles 8 and 11 of the EU Charter of Fundamental Rights. These reasons hinge on a premise. By quoting its case law (C-275/06, that is, *Promusicae*), the Court affirms that none of the rights to intellectual property are either "inviolable," or "absolute," but rather, they should be balanced against the protection of other fundamental rights (C-360/10, §§ 42-43). Therefore, on the one hand, "such an injunction [requiring the installation of the contested filtering system] would result in a serious infringement of the freedom of the hosting service

provider to conduct its business since it would require that hosting service provider to install a complicated, costly, permanent computer system at its own expense, which would also be contrary to the conditions laid down in Article 3(1) of Directive 2004/48, which requires that measures to ensure the respect of intellectual-property rights should not be unnecessarily complicated or costly" (*op. cit.*, § 46).

On the other hand, in the opinion of the Court, this sample of legal automation should be reckoned as illegitimate because indiscriminate. The installation of such filtering system would not only involve the identification, systematic analysis and processing of information connected with the profiles created on the social network by its users, hence impinging on how personal data shall be protected. "Moreover, that injunction [to install the filtering system] could potentially undermine freedom of information, since that system might not distinguish adequately between unlawful content and lawful content, with the result that its introduction could lead to the blocking of lawful communications" (*op. cit.*, § 50). Therefore, not only the kind of legal automation, which was at stake in *Netlog*, has to be deemed as illegitimate, in order to protect such basic rights as freedom to receive or impart information, or the protection of personal data, but it is noteworthy that no balancing was needed in the case. In the phrasing of the Court, the EU law "must be interpreted as precluding a national court from issuing an injunction against a hosting service provider which requires it to install a system for filtering," as the system described above in the previous section.

More recently, on 8 April 2014, a similar verdict was returned with regard to the 2006 EU data retention directive (joined cases C-293/12 and C-594/12). Justices in Luxembourg declared the latter invalid, because D-2006/24/EC infringed the proportionality principle by generally affecting "all persons and all means of electronic communication as well as all traffic data without any differentiation, limitation or exception being made in the light of the objective of fighting against serious crime" (*op. cit.*, § 56). In addition, no "objective criterion" was laid down in the directive so as to determine who could access and subsequently make use of the data retained in accordance with the directive 24 from 2006 and "what is strictly necessary in the light of the objective pursued" (*op. cit.*, § 62). As occurred in the *Netlog* case, no balancing was thus required to declare such provisions invalid and yet, dealing with the legitimacy of how far norms of legal automation can go, we should not leap to conclusions. Whilst, theoretically speaking, it is feasible to mull over either cheap filtering systems that do not end up "in a serious infringement of the freedom of the hosting service provider," or smarter self-enforcing technologies that adequately distinguish unlawful content from people's lawful communications, where we should legally draw the line between the pros and cons of legal automation appears even harder. Next section aims to explain why this is the case.

3.3 The Open Issues of Legal Automation

It is still unclear what type of legal automation would ultimately be legitimate in EU law. Three examples are fruitful to illustrate the point. First, some controversial provisions of the UK Digital Economy Act (DEA) from 2010, bring us back to uncertainties and dilemmas that end up with the preliminary ruling in the *Netlog* case. DEA lays down an "initial obligations code" that should impose on ISPs the duty to notify subscribers of copyright infringements reports received from copyright owners, and to provide copyright infringement lists to copyright owners, in addition to "technical obligations," some of which include a "technical obligations code." Certain ISPs, such as British Telecom, claimed that such provisions are illegitimate pursuant to EU directives on data protection, copyright, freedom to conduct a business, and so forth. However, two British courts endorsed the opinion of some powerful copyright-holders and simply ignored, or deemed as irrelevant, the jurisprudence of the EU Court of Justice. In the wording of the Court of Appeal in London, on 6 March 2012, "a certain amount of energy was expended before us on the recent judgement of the Court of Justice in *Scarlet*… which concerned the compatibility with the Privacy and Electronic Communications Directive and other directives of a court injunction against an ISP requiring it to install a system for filtering electronic communications in order to identify and block the transfer of files infringing copyright. Both the Advocate General and the Court referred to *Promusicae*, in terms that do not in my view cast any great light on that ruling; but I see nothing in

the case to support the limited scope that the applicants seek to give to the ruling in *Promusicae*" (CI/2011/1437, n. 82).

Second, as an outcome of the Snowden revelations, no reform has so far materialised on either side of the Atlantic. Whilst UK and France have legislated to extend surveillance powers, and the NSA Inquiry Committee of the German Bundestag was spied on by the CIA, the "USA Freedom" bill failed in America. Remarkably, in January 2012, the EU Commission presented a new *Police and Criminal Justice Data Protection Directive* that, pursuant to Articles 19(2) and 38 of the Proposal, should protect rights and freedoms of the individuals with appropriate technical and organizational measures, "which meet in particular the principles of data protection by design and data protection by default" (*op. cit.*, n. 38). Moreover, the Commission significantly referred to "the level of protection of the rights and freedoms of individuals with regard to the processing of personal data by competent authorities for the purposes of prevention, investigation, detection or prosecution of criminal offences," which "must be equivalent in all Member States" (Seventh Considerandum of the proposed directive). However, three years later, especially after the terrorist attacks in Paris from January 2015, governments have recommended an intensification of surveillance programs that suggests that such plans, as the 2012 Commission proposal, will be abandoned, at least, for quite a long time.

Third, even the EU Court of Justice has changed its mind with the ruling in *Google v. AEPD*, i.e. the famous case on the right to be forgotten from 13 May 2014 (C-131/12). Here, for overt political reasons, Justices in Luxembourg established that search engines, such as Google's, should be conceived as "data controllers" (*op. cit.*, §§ 33 and 34), thereby overruling what declared in the *Google v. Louis Vuitton case* on 23 March 2010. In this latter occasion, the opinion was that liability of online referencing service providers ultimately depends on "the actual terms on which the service is supplied." In other words, according to the judges in Luxembourg, it was necessary to determine, at least until 13 May 2014, "whether the role played by that service provider is neutral, in the sense that its conduct is merely technical, automatic and passive, pointing to a lack of knowledge or control of the data which it stores" (§ 114 of the decision). Some reckon that, by reversing this idea and claiming that search engines algorithms are no longer ontologically neutral, the Court anticipated what the Commission has proposed with Article 17 of the new EU data protection regulation from January 2012, namely a new set of duties and obligations for ISPs in the name of the right to be forgotten. Yet, in November 2013 and later, in March 2014, the EU Parliament passed a set of amendments that redesign this set of rules so much, that even the reference to the right to be forgotten has disappeared in the new text.

In light of these examples, what appears clear is the urgency of a normative standpoint with which we should tackle the challenges of legal automation and its hard cases. This requires intelligence and moreover, cannot be straightforwardly made subject to legal automation. Rather, what is at stake here concerns critical decisions vis-à-vis the insurance of fundamental legal rights, much as choices of technological dependence, delegation and trust, that have to ascertain the good mix between legal automation and public deliberation. Let us deepen this complex set of issues in the final part of this paper.

4 The Governance of Legal Automation

Regardless of the legal field under exam, the delegation of decisions to automated systems and hence, the governance of legal automation has to tackle two common magnitudes of complexity. Along with the aim to embed normative constraints into technology, as was illustrated above in Section 2.2, attention should be drawn to the interplay between law & technology and moreover, to the intent of the law to govern the process of technological innovation in such a way, that legal regulation should neither hinder the advance of technology, nor require over-frequent revision to tackle such a progress. This latter perspective on the regulative aims of the law has not to be confused with the field of techno-regulation, or legal regulation by design, which was presented in Section 2.3 and concerns how current advancements of technology have obliged legislators and policy makers to forge more sophisticated ways to think about legal enforcement. In the case of the aim of the law to regulate technological innovation, i.e. the law conceived as a "meta-technology" (Pagallo 2013), focus is on the different normative purposes that the law can have,

including that which scholars often dub as the "technological neutrality" of the law. For example, according to Chris Reed (2012), we should differentiate between (a) technological indifference, i.e. legal regulations which apply in identical ways, whatever the technology, such as the right to authorize communication of a work to the public in the field of copyright law; (b) implementation neutrality, so that regulations are by definition specific to that technology and yet, they do not favour one or more of its possible implementations, e.g. the signature of e-documents; and, (c) potential neutrality of the law that sets up a particular attribute of a technology, although lawmakers can draft the legal requirement so that even non-compliant implementations can be modified to become compliant. Alternatively, Bert-Jaap Koops has proposed to distinguish four main legislative purposes, such as: (a) the achievement of particular effects, e.g. preventing harm-generating behaviour from occurring, or decreasing its impact, through the means of legal automation; (b) functional equivalence between online and offline activities, e.g. security measures for atomic plants facilities and their IT systems; (c) non-discrimination between technologies with equivalent effects; and, (d) future-proofing of the law that should be compatible with the advance of technology, so as not to be often revised in order to keep the pace of such a progress (Koops 2006).

Yet, the different and even opposite ways in which we can grasp the normative purposes of the law as a meta-technology recommend to expand our view. First, attention should be drawn to a meta-regulatory approach to the field of legal automation, with which determine whether, or to what extent, lawmakers shall not (or cannot) delegate decisions to automated systems. Second, focus should be on the impact of technology on the formalism of the law, and how the latter competes with further regulatory systems. Third, we have to pay attention to the principles and values which are at stake with the delegation of decisions to automated systems, namely the institutional dimension of the law with matters of interpretation and deliberation. Fourth, the distinction between automatic and non-automatic decisions of the law, and their legitimacy, may entail a class of legal problems, i.e. the hard cases of the law, where disagreement can revolve around semantics, or legal reasoning, or the role and logic of the principles in the system. Each of these issues is deepened in the four parts of this section on the limits of legal automation (4.1); competing regulatory systems (4.2); the institutional dimension of the law (4.3); and its hard cases (4.4). Then, the time will be ripe for the conclusions of this paper.

4.1 The Limits of Legal Automation

The pros and cons of legal automation are hard to disengage. Consider first of all some intrinsic limits that affect the implementation process of the legal tasks delegated to automated systems and the ability of the law to anticipate the evolution of technology. As stressed above in Section 2.1, on the one hand, the delegation of decisions to automated systems does not cover every aspect of the law, nor all the legal solutions. Generally speaking, the difficulty concerns how to weld the syntactic levels of automation into the semantic dimension of the law. Here, we can speak about the "internal" limits of legal automation. On the other hand, we do not have to endorse any techno-deterministic stance to accept that which was mentioned in the previous section, namely, that legal systems are not always capable to predict and anticipate every technology change, so as to catch up with the race of science and technological innovation. Moreover, we insisted on how the intent of the law to govern this process should not hinder the advance of technology. Both the descriptive and normative aspects of this view suggest what we can sum up as the "external" limits of legal automation. They regard the limits of prediction and anticipation by the law, which restraints the set of legal issues that can be delegated to the decisions of an automated system.

As a result, the internal and external limits of legal automation cast light on the fact that we cannot draw a line between the pros and cons of legal automation in its own terms. We are in fact confronted with a dialectics, i.e. the interplay between law and technology, that cannot be solved like a Gordian knot, with a sword. Rather, a balance should be struck between automated systems and the traditional tools of the law, so as to determine whether a series of tasks that were usually carried out through such traditional means, i.e. the "ought to" of the law, can finally be entrusted to any automated system. Accordingly, in the basket of legal goods, we find a necessary and even inescapable mix of automation and non-automation that is not something entirely new. In the history of jurisprudence and the legal tradition, after all, we find the classical distinction between

an automatic interpretation and application of the law (e.g. Hart 1961), and vice versa an interpretation and application of the law which stems from meditation, criticism and prudent evaluation of the legal principles and rules of the system. This latter perspective suggests a meta-regulatory approach to the limits of legal automation, that has to take into account the regulative aim of the law as a system which competes with other regulatory systems and furthermore, as an institutional sphere in which we should strike the fair balance between automation and non-automation. Next section dwells on the first target of this meta-regulatory approach.

4.2 Competing Regulatory Systems

The interplay between law and technology can be understood as the interaction between competing regulatory systems that contend against further regulatory systems, as the forces of the market and of social norms. Every regulatory system claims to govern social interaction by its own means and with the pros and cons that we already examined in the previous sections. Such regulatory claims may either clash, or reinforce each other or finally, a regulatory system can render the claim of another regulatory system superfluous. Whatever the scenario we consider, however, such a competition does not take place in a normative vacuum but rather, is structured by the presence of values and principles. The normative contexts that we take into account, e.g. the pros and cons of legal automation, can thus be characterized by a shared set of values and principles, i.e. a general social agreement, or not. This bifurcation is critical, because it tells us something new about the process of legal automation from a meta-regulatory point of view. The issues brought on by the delegation of decisions to automated systems do not only depend on the degree of predictability and reliability of such automated decisions. In addition, these issues hinge on the degree of social agreement, or disagreement, that characterize the normative context with which we are dealing.

This normative stance draws the attention to another aspect of the problem that is often overlooked, or underestimated. Decisions delegated to machines do affect assets and interests that can be measured with the degree of "social acceptability" that concerns the risk inherent in the automation process. Consider for example human interaction with personal robots that may involve emotional, physical and physiological activities that have a cost even for adult human beings. Some wonder if it is "ethically justifiable to aim to create robots that people bond with, *e.g.*, in the case of elderly people or people with special needs" (Dautenhahn 2007: 699). Still, the technical and legal governance of how decisions delegated to machines may affect assets and human interests, does not entirely depend on the degree of social acceptability, but also, or above all, on the degree of social cohesion that concerns values and principles that are at stake with those assets and interests. Going back to the field of robotics, whether humans will get the same payoff and gratification from their interaction with such artificial agents, as they do with other human fellows, is a question that mostly depends on the cultural context (and the type of robotic application) with which we are confronted. Rather than measured in connection with levels of social acceptability, technological dependence and the corresponding grade of delegation and autonomy have thus to be comprehended in accordance with the degree of social cohesion that exists in the normative context in which the consequences of tasks and decisions delegated to automated systems, have to be evaluated. The stronger such a social cohesion is, the higher a risk in the automation process that can be socially accepted, that is, the normative assessment of not entirely predictable consequences of entrusted tasks and decisions to machines.

Against this backdrop, the next step of the analysis has to do with the institutional forum within which such normative assessment and the degree of social cohesion shall be measured.

4.3 The Institutional Dimension of the Law

The formation and stewardship of the formal and informal rules that shall govern the process of legal automation have to address the twofold set of issues mentioned above in the previous sections, namely: (i) the limits of legal automation and hence, how to strike a balance between delegation of decisions to automated systems and non-delegation; and, (ii) the normative context in which the consequences of such a balance will occur. An institutional forum is thus required, in order to attain the necessary legal and ethical framework for a public deliberation on the interplay between law and technology.

Drawing on the tradition of legal philosophy and jurisprudence, we can grasp the tension between automation and non-automation in terms of interpretation and application of the law. It would be a mistake to grasp this tension as entirely provoked by the evolution of ICTs and digital technologies and, in more general terms, by how current societies depend on technologies that concern information as their vital resource. After all, lawyers often deal with a complex set of concepts and notions that, still, leave no doubts as to how to apply them in the legal domain. Those are, in the words of Herbert Hart, the cases where legal issues are pretty "plain," that is, "where the general terms seem to need no interpretation and where the recognition of instances seems unproblematic or 'automatic'… where there is general agreement in judgements as to the applicability of the classifying terms" (Hart 1961: 121).

However, this stance should be furthered since it generally applies to all human things. The distinction between automatic and non-automatic decisions is not coextensive with the distinction between humans and machines. Rather, it is inherent to the nature of human beings (Kahneman 2011). Many of our decisions are not the result of meditation, criticism and prudent evaluation, but of the automatic and reiterated application of already acquired competences. Individuals often need to decide thoughtlessly, because of the particular circumstances of the case, for lack of time, or of information. In addition, human behavior frequently is not guided by conscious choices and deliberation but by the need to adapt to the environment. Not all human activities require intelligence and *pour cause*. The idea that the notion of legal automation is always, or should principally be, framed in terms of human and artificial intelligence is hence a largely abused posture that perhaps, should rather be traced back to the notion of delegated tasks, that is, to the conceptual area of trust (Durante 2010; Pagallo 2010).

Going back to the field of legal interpretation, which includes formalism and realism, dogmatism and skepticism, and so forth, the distinction between automatic and non-automatic decisions relies on the distinction between the aforementioned plain cases of jurisprudence, and its "hard cases." The distinction has been the subject of a much lively debate in the philosophy of law. Of course, there is no room to reconstruct here the debate nor would it serve the purpose of this paper. Rather, suffice it to refer to some of the main existing literature (Hart 1961; Dworkin 1977, 1986; Shapiro 2007) and to briefly restate the meaning of the distinction in the next section, through the words of its most celebrated interpreter, Herbert Hart. In light of this distinction between plain and hard cases of the law, we return to matters of legal automation and its governance in the conclusions of the paper. The aim is to ascertain whether the balance that should be struck between delegation of decisions to automated systems and non-delegation, much as the normative context in which the consequences of such a balance will occur, spark disagreement in the community and how we should react before such cases.

4.4 Between Plain and Hard Cases

Hart's parallel between legal plain cases and automation, as quoted above in the previous section, draws the attention to three key aspects of the current debate on whether, or to what extent, legal systems should delegate decisions to automated systems. First, these cases need no particular recourse to human intelligence for their interpretation and application, so much so, that they appear straightforward and open to automation. Second, such cases are not uncommon but rather, they are the familiar ones, i.e. those which constantly recur in similar contexts since the law is not confronted with something radically new or problematic. Third, the automatic nature of the legal plain cases follows a generally shared opinion on the terms of their application. More extensively and less formalistically, we can say that what makes a legal case "plain" is a commonly shared (and sufficiently clear) connection between the legal decision of the case and its normative context, as introduced above in Section 4.2.

The existence of a conceptual and axiological common core, moreover, is not an accidental feature but the very condition of possibility of the law, understood as an instrument of social control. As properly stressed by Brian Bix, "Hart was concerned with the problem of social control through law: not questions of strategy or political theory, of how social control could best be effected, but the preliminary question of how social control could even be possible. How can a government guide its population's actions on the basis of legislation and precedent, and to what extent will those means necessarily need supplementation? Hart stated: 'If it were not possible to

communicate general standards of conduct, which multitudes of individuals could understand, without further direction, as requiring from them certain conduct when occasion arose, nothing that we now recognize as law could exist"' (Bix 1991: 54; citing Hart 1961: 121). On this basis, it is not hard to tell what a legal hard case is. Here, contrary to the plain cases, we deal with general disagreement that may regard: (a) the meaning of the terms framing the legal question; (b) the ways such terms are related to each other in legal reasoning; or, (c) the role of the principles that are at stake in the case.

This sequence, from legal terms to principles, makes clear that general disagreement may not only hinge on the interpretation of the legal texts but moreover, on different values and principles of the normative context, within which the issue under exam is appraised. This latter scenario seems to trigger a vicious circle, much as "the chicken or the egg" causality dilemma. The law is here confronted with something new and problematic that, on the one hand, makes a merely automatic application of the law insufficient and yet, on the other hand, needs meditation, criticism and a prudent evaluation that, in Hart's words, entails a "further direction" that draws on the background of values and principles that constitute the legal framework of the hard legal decision. As an illustration of this stalemate, consider current debate on the use of autonomous lethal weapons on the battlefield, matters of liability for contracts executed by robots and other artificial agents, or forms of distributed responsibility for complex multi-agent systems, that were stressed in the introduction of this paper. The apparent circularity of the legal hard cases is however misleading, once we distinguish the different role that the agreement of the plain cases plays vis-à-vis the disagreement of the hard cases in the legal domain. In the first case, that is, general agreement that makes the parallel between plain cases and automation feasible, such an agreement represents the condition for the existence and normal functioning of the law, through standards of conduct, such as norms, values, and principles, that need no "further direction." Vice versa, the different types of disagreement that make a legal case hard, illustrate what the law is, namely the concept of law. Such cases highlight all the relevant standards of conduct, i.e. norms, values, or principles, that can be adopted as the basis of a legal decision and yet, require a supplement of direction in terms of human intelligence.

The implementation of legal automation, as a matter of principle, can thus concern the conditions of existence and normal functioning of the law: in other words, they do not automatically affect the appraisal of what the law is, i.e. the legally relevant standards of conduct. This scenario has been illustrated in the first part of the paper with work in such fields, as legal ontologies and more generally, AI & the law: see above in Section 2.

On the other hand, the hard cases play a crucial role, for they delimit the process of legal automation and moreover, require an institutional space of interpretation in which the legal texts are understood and evaluated within the normative context of the law. Whether the hard cases of the law should be addressed through "reasonable compromises" (Hart), or alternatively, by the means of a morally coherent interpretation that best fits "the integrity of the law" (Dworkin), is a meta-hard case of jurisprudence that we can leave aside in this context. Rather, focus should be on a new problematic set of legal and ethical challenges that bring us back to the values and principles structuring the normative context of the law. This process of interpretation sets the framework for a public discussion and deliberation that aims to address the relation between law and technology innovation. The weaker the degree of social cohesion that exists in the normative context, in which the consequences of tasks and decisions delegated to automated systems have to be evaluated, the lower the risk in the automation process that will be socially accepted for the unpredictable consequences that may follow the set of entrusted tasks and decisions to machines.

5 Conclusions

The paper has examined the pros of legal automation (Section 2), and its cons (Section 3), in order to flesh out goals and values that are at stake with choices of technological dependence, delegation and trust, and determine the good mix between legal automation and public deliberation (Section 4). On the one hand, the limits of legal automation do not hinge on any pretended semantic irreducibility of human decisions to automated outputs. These limits have to be understood against the backdrop of two different classes of legal decisions. The delegation of decisions to automated

systems does not automatically affect the relevant standards of conduct which the law takes into account, since the implementation of legal automation can concern the conditions of existence and normal functioning of rules, values, and principles, that substantiate the normative context of the law. This is set of the legal plain cases that was illustrated with the pros of legal automation.

On the other hand, there is a further class of legal decisions, that is the hard cases of the law, that should not be entrusted to automated machines, whether or not this is technically feasible. As the benchmark of our current ethical and legal challenges, such hard cases require meditation, criticism and a prudent evaluation of the principles and rules of the system. In addition, this process of interpretation has to be comprehended within a framework for public discussion and deliberation on the values and principles that structure the normative context of the law. As stressed time and again throughout this paper, today's innovation and evolution in automation is triggering an increasing number of legal hard cases, for they confront the law with something radically new and problematic, such as the new surveillance society, the new scenarios of cyber warfare and automatic lethal machines, down to the fact that, for the first time ever, human societies depend on information as their vital resource. No settled values and principles guide the normative context of the assets and interests affected by the legal decisions potentially delegated to machines in such cases. Here, the "social acceptability" of the risk inherent to the automation process is controversial and furthermore, this controversy is rooted into a deeper general disagreement, or lack of "social cohesion," which regards the values and principles that form the normative context of every legal decision.

Correspondingly, the key question does not revolve around whether or not there is an irreducible semantic core in the act of deciding, which should thus be entrusted only to human beings. Multiple levels of semantic and axiological complexity exist according to different classes of legal decisions and a number of intricate cognitive tasks, as a matter of fact, has already been delegated to machines. Hence, the question is not how far the process of legal automation can go but rather, whether the distinction between plain and hard cases can be subjected to a process of legal automation. We reckon that this distinction cannot be entrusted to machines but should be reserved to human beings that still bear full responsibility for the judgment of what is socially, ethically, and legally "plain" and "hard" in human affairs. The line between the pros and cons of legal automation cannot be drawn in its own terms.

References

Abou-Tair, D. el Diehn I., and Stefan Berlik (2006) An Ontology-based Approach for Managing and Maintaining Privacy in Information Systems, *Lectures Notes in Computer Science* 4275: 983-994

Bix, Brian (1991) H.L.A. Hart and the 'open texture' of language, *Law and Philosophy*, 10:1, 51-72

Breuker, Joost et al. (eds.) (2009) *Law, Ontologies and the Semantic Web. Channelling the Legal Information Flood.* Amsterdam: IOS Press

Casanovas, Pompeu, Pagallo, Ugo, Palmirani, Monica and Giovanni Sartor (eds.) (2014) *AI Approaches to the Complexity of Legal Systems. Law, Social Intelligence, nMAS and the Semantic* Web. Berlin-Heidelberg: Springer

Casanovas, Pompeu, Pagallo, Ugo, Sartor, Giovanni and Gianmaria Ajani (eds.) (2010) *AI Approaches to the Complexity of Legal Systems. Complex Systems, the Semantic Web, Ontologies, Argumentation, and Dialogue*. Berlin-Heidelberg: Springer,

Casellas, Núria (2011) *Legal Ontology Engineering. Methodologies, Modelling Trends, and the Ontology of Professional Judicial Knowledge*. Berlin-Heidelberg: Springer

Cavoukian, Ann (2010) Privacy by Design: The Definitive Workshop, *Identity in the Information Society*, 3(2): 247-251

Durante, Massimo (2010) What Is the Model of Trust for Multi-agent Systems? Whether or Not E-Trust Applies to Autonomous Agents, *Knowledge, Technology & Policy*, 23(3-4): 347-366

Durante, Massimo (2011a) Normativity, Constructionism, and Constraining Affordances, in *Ethics and Politics*, XIII (2): 180-200

Durante, Massimo (2011b) Rethinking Human Identity in the Age of Autonomic Computing: the Philosophical Idea of Trace, in M. Hildebrandt - A. Rouvroy (eds.), *Law, Human Agency and Autonomic Computing. The Philosophy of Law meets the Philosophy of Technology*, London–New York,

Routledge, 85-103
Durante, Massimo (2013) Dealing with Legal Conflicts in the Information Society. An Informational Understanding of Balancing Competing Interests, in *Philosophy & Technology*, 26(4): 437-457
Dworkin, Ronald (1977) *Taking Rights Seriously*. Cambridge, MA: Harvard University Press
Dworkin, Ronald (1986), *Law's Empire*. Cambridge, MA: Harvard University Press
Flanagan, Mary, Howe, Daniel C., and Helen Nissenbaum (2008) Embodying Values in Technology: Theory and Practice. In *Information Technology and Moral Philosophy*, J. van den Hoven and J. Weckert (eds.), 322-353. Cambridge University Press, New York
Froomkin, Michael A. (2000) The death of privacy? *Stanford Law Review*, 498: 1461-1543
Hart, Herbert (1961) *The Concept of Law*. Oxford: Oxford University Press
Iaconesi, Salvatore (2014) Ubiquitous Commons: A Browser Plugin to Enable User-generated, Peer-to-peer-enacted Licensing Models for Online Data, available at http://www.ubiquitouscommons.org (last accessed 1 June 2015)
Kelsen, Hans (1949) *General Theory of the Law and the State*, trans. A. Wedberg. Cambridge, Mass.: Harvard University Press
Koops, Bert-Jaap and Ronald Leenes (2014) Privacy Regulation Cannot Be Hardcoded: A Critical Comment on the "Privacy by Design" Provision in Data Protection Law, *International Review of Law, Computers & Technology*, 28: 159-171
Leenes, Ronald and Federica Lucivero (2014) Laws on Robots, Laws by Robots, Laws in Robots: Regulating Robot Behaviour by Design, *Law, Innovation and Technology*, 6(2): 193-220
Lessig, Lawrence (2004) *Free Culture: The Nature and Future of Creativity*. Penguin, New York
Lioudakis, Georgios, Koutsoloukasa, Eleftherios, Tselikasa, Nikolaos, Kapellakia, Sofia, Prezerakosa, Georg, Kaklamani, Dimitra and Iakovos Venieris (2007) A Middleware Architecture for Privacy Protection, *The International Journal of Computer and Telecommunications Networking*, 51(16): 4679-4696
Lyon, David (1994) *The Electronic Eye: The Rise of Surveillance Society*. University of Minnesota Press
Marx, Gary (2002) What's New About the "New Surveillance"? Classifying for Change and Continuity, *Surveillance & Society*, 1(1): 9-29
Mitre, Hugo, González-Tablas, Ana Isabel, Ramos, Benjamin and Arturo Ribagorda (2006) A Legal Ontology to Support Privacy Preservation in Location-based Services, *Lectures Notes in Computer Science*, 4278: 1755-1764
Pagallo, Ugo (2006) *Introduzione alla filosofia digitale: da Leibniz a Chaitin*. Torino: Giappichelli
Pagallo, Ugo (2011a) Designing Data Protection Safeguards Ethically, *Information*, 2(2): 247-265
Pagallo, Ugo (2011b) Online Security and the Protection of Civil Rights: A Legal Overview, *Philosophy & Technology*, 26(4): 381-395
Pagallo, Ugo (2012) Cracking down on Autonomy: Three Challenges to Design in IT Law, *Ethics and Information Technology*, 14(4): 319-328
Palmirani, Monica, Pagallo, Ugo, Casanovas, Pompeu and Giovanni Sartor (eds.) (2012) *AI Approaches to the Complexity of Legal Systems. Models and Ethical Challenges for Legal Systems, Legal Language and Legal Ontologies, Argumentation and Software Agents*. Berlin-Heidelberg: Springer
RoboLaw (2014) Guidelines on Regulating Robotics. EU Project on Regulating Emerging Robotic Technologies in Europe: Robotics facing Law and Ethics, September 22
Sartor, Giovanni, Palmirani, Monica, Francesconi, Enrico and Maria Angela Biasotti (2011) *Legislative XML for the Semantic Web. Principles, Models, Standards for Document Management*. Dordrecht: Springer
Shapiro, Scott J. (2007) The 'Hart-Dworkin' debate: a short guide for the perplexed, *Public Law and Legal Theory Working Paper Series*, 77, Michigan Law School
Simon, Herbert A. (1996) *The Sciences of the Artificial*. Cambridge, Mass., MIT Press
WP 29 (2009) *The Future of Privacy*. EU Working Party art. 29 D-95/46/EC: WP 168, December 1st
Yeung, Karen (2007) Towards an Understanding of Regulation by Design. In *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes*, R. Brownsword and K. Yeung (eds.), 79-108. Hart, London
Zittrain, Jonathan (2007) Perfect Enforcement on Tomorrow's Internet. In *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes*, R. Brownsword and K. Yeung (eds.), 125-156. Hart, London