The Pros and Cons of Legal Automation and its Governance

*Ugo Pagallo and Massimo Durante**

The article examines the field 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 human behaviour through codes, IT architectures, and design. By stressing both benefits and shortcomings of legal automation, the article does not suggest that the latter is something "neutral". Rather, making legal reasoning and enforcement automatic, so that even a machine can process and understand this information, should be conceived as a set of constraints and affordances that transform, or reshape, the environment of people's interaction and moreover, 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.

I. Introduction

This article examines the field 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. Admittedly, "legal automation" is a broad notion: in this context, we refer to it as the grandfather of current research in AI & the law, the German philosopher G.W. Leibniz, did three and a half centuries ago. Leibniz's aim was to turn legal arguments into computing through combinatorial analysis, probability calculus, and binary arithmetic.[1] Even a machine, at the end of the day, should process and understand this kind of information. What consequences follow in the legal domain?

The article is presented in three parts. Section II illustrates the pros of legal automation as stressed time and again by experts of such fields as legal informatics, AI and the law, and more. The intent of this research is to improve the efficiency, consistency, comprehensibility, and predictability of legal and judicial systems through machine learning and data mining techniques for legal applications, computational models of legal reasoning, ICT applications to support the legal domain, and so forth. Section III examines the cons of legal automation, i.e. the bread and butter of work on the new surveillance society, the death of privacy, legal regulation by design, codes, and IT architectures. This kind of debate sheds light on the implicit values of technology and its invisibility, the challenges of self-enforcing technologies, down to the lack of public debate brought on by automation, as occurs for example with the use of drones on the battlefield with no parliamentary authorization. The emphasis on both advantages and drawbacks of legal automation does not intend to suggest that the latter is something "neutral," namely a simple means to achieve whatsoever end. Rather, the aim to make legal reasoning and enforcement automatic, should be conceived as a set of constraints and affordances,[2] 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 as-

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

[1] Ugo Pagallo, *Introduzione alla filosofia digitale: da Leibniz a Chaitin* (Torino: Giappichelli, 2006).

[2] Massimo Durante, "Normativity, Constructionism, and Constraining Affordances", XIII (2) *Ethics and Politics* (2011), pp. 180 *et sqq.*; Ugo Pagallo, "Good Onlife Governance: On Law, Spontaneous Orders, and Design", in Luciano Floridi (ed.), *The Onlife Manifesto: Being Human in a Hyperconnected Era* (Dordrecht: Springer, 2015), pp. 161 *et sqq*.

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

pects of legal representation and resolution: think about the right of the individuals to have a say in the decisions affecting them on the internet.

Against this backdrop, the final part of the article draws the attention to the formation and stewardship of the formal and informal rules that regulate such a crucial aspect of current legal systems as their automation (section IV). The governance of legal automation is deepened through the analysis of both the internal and external limits (section IV.1), the theoretical framework of the interaction between law and technology as competing regulatory systems (section IV.2), and the need for an institutional forum for deliberation (section IV.3), all of which end up with the distinction between plain and hard cases of the law (section IV.4). Since cases of disagreement brought on by legal automation will probably be the main subject of lawyers, scholars, and policy makers for quite a long time, the overall aim of the article is to provide a normative stance with which to tackle such legal hard cases and determine the good mix between legal automation and public deliberation.

II. The Pros of Legal Automation

Three and a half centuries after Leibniz's work, his dream to make legal reasoning and enforcement automatic 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.[3] 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 in AI & the law, namely legal ontologies. Once familiar with this field (section II.1), and how it works in the legal domain of data protection (section II.2), limits and risks of this implementation should be clear. At that point we'll be ready for the other side of the coin: the cons of legal automation.

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.[4] The objective is that even a machine should comprehend and process this very information, by 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. However, many 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, or prohibition, but context-dependent legal notions, e.g. personal data and security measures, as well. This difficulty does not mean 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 ex-

[3] Pompeu Casanovas, Ugo Pagallo, Monica Palmirani and Giovanni Sartor (eds.), *AI Approaches to the Complexity of Legal Systems. Law, Social Intelligence, nMAS and the Semantic Web*. (Berlin-Heidelberg: Springer, 2014); Pompeu Casanovas, Ugo Pagallo, Giovanni Sartor, and Gianmaria Ajani (eds.), *AI Approaches to the Complexity of Legal Systems. Complex Systems, the Semantic Web, Ontologies, Argumentation, and Dialogue* (Berlin-Heidelberg: Springer, 2010); Monica Palmirani, Ugo Pagallo, Pompeu Casanovas et al. (eds.), *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, 2012).

[4] Joost Breuker, Pompeu Casanovas, Michel C.A. Klein et al. (eds.), *Law, Ontologies and the Semantic Web. Channelling the Legal Information Flood* (Amsterdam: IOS Press, 2009); Casanovas, *Complex Systems, supra* note 3.

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

amines 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"[5].

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.[6] Likewise, consider work on legal ontologies for the support of privacy preservation in location-based services, the management of information systems, or middleware architectures for data protection, each of which aims at integrating smaller parts and sub-solutions into the design of the project.[7] 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.[8] The overall idea is to embed legal constraints, or safeguards, into information systems and other technologies, 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. Two Roads to Automatic Privacy by Design

The idea of embedding privacy safeguards into information systems and other technologies is nothing new. The Ontario's Privacy Commissioner, Ann Cavoukian, invented the formula "privacy by design" in the late 1990s. 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."

The idea of embedding privacy safeguards into technology may or may not entail any automation. Thus, we have to 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.[9] 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"[10]. 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 reme-

[5] Herbert A. Simon, *The Sciences of the Artificial*. Cambridge (Mass., MIT Press, 1996), at p. 128.

[6] Giovanni Sartor, Monica Palmirani, Enrico Francesconi, et al., *Legislative XML for the Semantic Web. Principles, Models, Standards for Document Management* (Dordrecht: Springer, 2011).

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

[8] Ugo Pagallo, "Designing Data Protection Safeguards Ethically", 2(2) *Information* (2011a), pp. 247 *et sqq*.

[9] Ronald Leenes and Federica Lucivero, "Laws on Robots, Laws by Robots, Laws in Robots: Regulating Robot Behaviour by Design", 6(2) *Law, Innovation and Technology* (2014), pp. 193 *et sqq*.

[10] RoboLaw, Guidelines on Regulating Robotics. EU Project on Regulating Emerging Robotic Technologies in Europe: Robotics facing Law and Ethics (22 September 2014).

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

dial. 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.[11] By making 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). 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. All these questions and issues introduce the second part of the article on the cons of legal automation.

III. The Cons of Legal Automation

There is a number of ethical and technical reasons why making legal protection automatic is problematic. As to the 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"[12]. In the case of data protection, introduced above in the previous section, 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. All in all, 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 mentioned above in section II.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 is unilaterally determined on the basis of technology.[13] 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"[14]. 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 warns, 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"[15].

Finally, the technical difficulties of achieving such a total control through design should be mentioned.

[11] Ann Cavoukian, "Privacy by Design: The Definitive Workshop", 3(2) *Identity in the Information Society* (2010), pp. 247 *et sqq*.

[12] Mary Flanagan, Daniel C. Howe and Helen Nissenbaum, "Embodying Values in Technology: Theory and Practice", in Jeroen van den Hoven and John Weckert (eds.), *Information Technology and Moral Philosophy* (Cambridge University Press, New York, 2008), pp. 322 *et sqq*.

[13] Ugo Pagallo, "Cracking down on Autonomy: Three Challenges to Design in IT Law" 14(4) *Ethics and Information Technology* 2012), pp. 319 *et sqq*.

[14] Jonathan Zittrain, "Perfect Enforcement on Tomorrow's Internet" in Roger Brownsword and Karen Yeung (eds.), *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes*, (London: Hart, 2007), pp. 125 *et sqq*.

[15] Lawrence Lessig, *Free Culture: The Nature and Future of Creativity* (Penguin: New York, 2004), p. 152.

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

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"[16]. As stressed above in section II.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"[17].

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). First of all, legal automation impacts on 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"[18]. By making legal enforcement automatic, the law is converted into a set of effects (B) that 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 cons of legal automation can be deepened with a particular case regarding the use of filtering systems on the internet (section III.1). This stance allows us to grasp why, at least in the EU, some of such filtering systems should be deemed as illegal (section III.2); and yet, even after the ruling of the Court of Justice in 2012, such a use is fated to remain an open issue (section III.3). After this analysis, we will be ready to examine the governance of legal automation (section IV).

1. Filtering Information on the Internet

A lively debate over what role internet intermediaries, or service providers ("ISPs"), should have, in order to ensure online security and the protection of individual rights, has occurred in Europe over the past years.[19] 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, or freedom to conduct a business. 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 from 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:

[16] Karen Yeung, "Towards an Understanding of Regulation by Design", in Brownsword, *Regulating Technologies*, *supra* note 14, pp. 79 *et sqq*.

[17] Bert-Jaap Koops and Ronald Leenes, "Privacy Regulation Cannot Be Hardcoded: A Critical Comment on the "Privacy by Design" Provision in Data Protection Law", 28 *International Review of Law, Computers & Technology* (2014), pp. 159 *et sqq*.

[18] Hans Kelsen, *General Theory of the Law and the State*, trans. A. Wedberg (Cambridge, Mass.: Harvard University Press, 1949).

[19] Ugo Pagallo, "Online Security and the Protection of Civil Rights: A Legal Overview", 26(4) *Philosophy & Technology* (2011b), pp. 381 *et sqq*.

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

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 to find out whether introducing a system for filtering that information and preventing files being made available which infringe copyright is lawful in the EU. In addition, the Belgian court asked whether there was 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.

2. Matters of Unbalance

There are two reasons why the Court of Luxembourg ruled that the filtering system discussed in *Netlog* was 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). Consequently, not only the kind of legal automation, 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."

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 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 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 required to declare such provisions invalid. Still, dealing with the legitimacy of how far norms of legal automation can go, we should not leap to conclusions. Whilst it

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

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. The next section aims to explain why this is the case.

3. The Open Issues of Legal Automation

It is still unclear what type of legal automation would ultimately be legitimate in EU law. Two 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 infringement 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 law. However, two British courts endorsed the opinion of some powerful copyright-holders and simply ignored 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, 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 had been 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 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 that even the reference to the right to be forgotten has been put in brackets 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 making legal reasoning and enforcement automatic. 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 safeguard of fundamental legal rights, much as choices of technological dependence and delegation, 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 the article.

IV. The Governance of Legal Automation

The delegation of decisions to automated systems has to tackle a twofold magnitude of complexity. Along with the aim to embed normative constraints into technology, the 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

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

progress. This latter perspective on the regulative aims of the law has not to be confused with techno-regulation, i.e. 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 law regulating technological innovation, i.e. the law conceived as a "meta-technology"[20], the 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[21], 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 an innovation.[22]

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. We propose four steps of analysis. First, a meta-regulatory approach to the field of legal automation should allow us to determine whether, and to what extent, lawmakers shall not (or cannot) delegate decisions to automated systems. Second, focus should be on the impact of technology on the formalisms 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 (IV.1); competing regulatory systems (IV.2); the institutional dimension of the law (IV.3); and its hard cases (IV.4). Then, the time will be ripe for the conclusions of this article.

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. The delegation of decisions to automated systems covers neither every aspect of the law, nor all legal solutions. The difficulty concerns how to weld the syntactic levels of automation into the semantic dimension of the law. Hence, 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 pace of science and technological innovation. We also 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 latter view suggest that which we can sum up as the "external" limits of legal automation. They regard the limits of prediction and anticipation that affect the law and restrain the number of legal issues that can be delegated to the decisions of an automated system.

[20] Ugo Pagallo, *The Laws of Robots: Contracts, Crimes, and Torts* (Dordrecht: Springer, 2013).

[21] Chris Reed, *Making Laws for Cyberspace* (Oxford: Oxford University Press, 2012), pp. 82 *et sqq*.

[22] Koops, "Privacy Regulation Cannot Be Hardcoded", *supra* note 17.

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

The internal and external limits of making legal reasoning and the functioning of the law automatic 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 traditional means, i.e. the "ought to" of the law, can finally be entrusted to an automated system. In the basket of legal goods, we find a necessary and even inescapable mix of automation and non-automation that is not 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 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 have to strike the fair balance between automation and non-automation. The next section dwells on the first target of this meta-regulatory approach.

2. Competing Regulatory Systems: From Social Acceptability to Cohesion

The interplay between law and technology can be grasped as the interaction between competing regulatory systems that not only may contend against each other, but also 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 stressed in the previous sections. Such regulatory claims may not only clash, but reinforce each other. In addition, a regulatory system can render the claim of another regulatory system superfluous. Whatever the scenario we consider, such a competition does not take place in a normative vacuum but rather is structured by the presence of values and principles.[23] The normative contexts that are taken into account, e.g. in connection with the pros and cons of legal automation, are thus characterized by either 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 standpoint. 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. Rather, these issues hinge on the degree of social agreement, or disagreement, that characterize the normative context under examination.

This normative stance draws the attention to another aspect of the problem that is often underestimated. Decisions delegated to machines, smart applications and even autonomous artificial agents, affect assets and interests that can be measured with the degree of "social acceptability" concerning the risk inherent in the automation process. Consider 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"[24]. Still, the technical and legal governance of how decisions delegated to smart machines and artificial agents may affect assets and human interests does not entirely depend on the degree of social acceptability but, more importantly, on the degree of social cohesion that concerns the 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 an open question that mostly depends on the cultural context (and the type of robotic application) with which we are confronted. Rather than simply measured by the levels of social acceptability, technological dependence and the corresponding grade of delegation and autonomy have thus to be comprehended in accordance with the set of values and princi-

[23] Massimo Durante, "Dealing with Legal Conflicts in the Information Society. An Informational Understanding of Balancing Competing Interests", 26(4) *Philosophy & Technology* (2013), pp. 437 *et sqq*.

[24] Kerstin Dautenhahn, "Socially intelligent robots: dimensions of human–robot interaction", 362(1480) *Philosophical Transactions* (2007), at p. 699.

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

ples that exist in the normative context in which the consequences of tasks and decisions delegated to automated systems are evaluated. The stronger the social cohesion, the higher the risk in the automation process that can be socially accepted through the normative assessment of not fully predictable consequences of tasks and decisions entrusted to machines and artificial agents.

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

3. The Institutional Dimension of the Law

The formation and stewardship of the formal and informal rules that govern the process of legal automation have to address the twofold set of issues mentioned above in the previous sections, namely: (i) how to strike a balance between delegation of decisions to automated systems and non-delegation; and, (ii) how to evaluate 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 any public deliberation on how law and technology should interplay.

This institutional dimension of the law can, of course, be understood in manifold ways. Let us draw here on the tradition of legal philosophy and jurisprudence. We can correspondingly conceive the tension between automation and non-automation as a matter of interpretation and application of the law. This latter perspective seems particularly fruitful, because it prevents the mistake to grasp the tension between automation and non-automation as entirely provoked by the evolution of ICTs and digital technologies and, in more general terms, by how current societies depend on information as their vital resource. Since ancient Roman times, lawyers have often dealt with a complex set of notions that, nevertheless, leave no doubts as to how to apply them in the legal domain. These are, in the words of Herbert Hart, the cases where legal issues are "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"[25]. The plain cases of the law abound in everyday life: buying a newspaper, shopping at a mall, or on the internet, enjoying a dinner at the restaurant, etc.

However, this stance should be widened. By insisting, time and again, on the limits of legal automation, it should be stressed that the distinction between non-automatic and automatic decisions is not coextensive with the difference between humans and machines. Rather, the distinction is inherent to the nature of human beings.[26] 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, lack of time, or of information. In addition, human behaviour 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*.

Going back to the field of legal interpretation, the distinction between automatic and non-automatic decisions relies on the difference between the plain cases of jurisprudence and its "hard cases." This distinction has been the subject of much lively debate in legal philosophy and of course, there is no room to reconstruct the whole debate nor would it serve the purposes of this article. Suffice it to refer to some of the main literature[27]; and furthermore, to briefly restate it through the words of its most celebrated interpreter, Herbert Hart. We do not have to buy his distinction between plain and hard cases of the law, to admit that Hart's parallel between legal plain cases and automation draws the attention to a key aspect of the current debate on whether, and to what extent, legal systems should delegate decisions to automated systems. That which has to be scrutinized concerns whether delegation of decisions to automated systems sparks disagreement in the community and hence, within the normative context in which

[25] Herbert Hart, *The Concept of Law* (Oxford: Oxford University Press, 1961), at p. 121.

[26] Daniel Kahneman, *Thinking Fast and Slow* (New York: MacMillan, 2011).

[27] Hart, *The Concept of Law, supra* note 26; Ronald Dworkin, *Taking Rights Seriously* (Cambridge, MA: Harvard University Press, 1977); Scott J. Shapiro, "The 'Hart-Dworkin' debate: a short guide for the perplexed", 77 *Public Law and Legal Theory Working Paper Series* (2007) (Michigan Law School).

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

the consequences of such decisions will occur. How should we react before such legal hard cases?

4. Between Plain and Hard Cases

We mentioned that the plain cases of the law need no particular recourse to human intelligence for their interpretation, because they appear straightforward and open to automation. Such cases are not uncommon but rather, they are the familiar ones, i.e. those which constantly recur in similar contexts since the law and people's interaction are not confronted with something radically new or problematic. In addition, 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 output of the case and its normative context.

Contrary to the plain cases, a legal hard case concerns 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 legal texts but, on different values and principles of the normative context under examination. This latter scenario seems to trigger a vicious circle, much as "the chicken or the egg" causality dilemma. The law is confronted with something new and problematic that, on the one hand, makes a merely automatic application of the law insufficient. On the other hand, such a case needs meditation, criticism and a prudent evaluation that refers to the values and principles that constitute the legal framework and yet, these principles and values are a source of disagreement. The apparent circularity of the legal hard case is however misleading, once we distinguish the different role that the agreement of the plain case plays vis-à-vis the disagreement of the hard case 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 as norms, values, and principles, that need no "further direction"[28]. The implementation of legal automation, as a matter of principle, can thus go hand-in-hand with the conditions of existence and normal functioning of the law, since this implementation does not automatically affect the legally relevant standards of conduct.

On the other hand, 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 nevertheless, require a supplement of direction in terms of human intelligence. The hard cases play hence a crucial role, for they delimit the process of legal automation and moreover, need 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. We prefer to stress 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. The set of legal hard cases brought on by automation, e.g. filtering systems on the internet and delegation of tasks to increasingly smart robots, claims a public discussion and deliberation that shall address the interplay between law and technology. 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 of the automation process that will be socially acceptable, i.e. the unpredictable consequences that may follow the set of entrusted tasks and decisions to machines and artificial agents. In light of this correlation, what should our normative stance be?

V. Conclusions

The article examined the pros of legal automation (section II), and its cons (section III), in order to flesh out goals and values that are at stake with choices of

[28] Hart, *The Concept of Law*, *supra* note 26.

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press

technological dependence, delegation and trust.[29] The aim has been to determine the good mix between legal automation and public deliberation, vis-à-vis matters of social acceptability and cohesion (Section IV). Drawing on this analysis, the conclusion is twofold: on the one hand, the limits of legal automation do not hinge on any pretended semantic irreducibility of human decisions to automated outputs. After all, the delegation of decisions to automated systems, as such, does not affect the relevant standards of conduct which the law takes into account. On the contrary, the implementation of legal automation can fit like hand to glove with the conditions of existence and normal functioning of rules, values, and principles, that substantiate the normative context of the law. This is the set of legal plain cases 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. The hard cases of the law require human understanding and interpretation as well as meditation, criticism and a prudent evaluation of the principles and rules of the system. Furthermore, this process of interpretation, understanding, and meditation, 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 article, 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, up 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 and artificial agents in such hard cases. Here, the "social acceptability" of the risk inherent to the automation process is rather controversial and what is more, this controversy is rooted into a deeper general disagreement, i.e. a lack of "social cohesion," regarding values and principles that form the normative context of every legal decision.

Correspondingly, the key question does not revolve around whether or not an irreducible semantic core exists in the act of deciding, which should thus be entrusted only to human beings. Multiple levels of semantic and axiological complexity do exist according to different classes of legal decisions, so that, as a matter of fact, a number of intricate cognitive tasks has already been delegated to machines and artificial agents. 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 and smart artificial agents 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.

[29] Massimo Durante, "What Is the Model of Trust for Multi-agent Systems? Whether or Not E-Trust Applies to Autonomous Agents", 23(3-4) *Knowledge, Technology & Policy* (2010), pp. 347 *et sqq*.

https://doi.org/10.1017/S1867299X00005742 Published online by Cambridge University Press