Algo-Rhythms and the Beat of the Legal Drum

(Article begins on next page)

24 November 2024

Algo-Rhythms and the Beat of the Legal Drum

1. Introduction

There is a panoply of algorithms 'out there.' They include also but not only combinatorial algorithms, such as graph and sequence algorithms, operating systems and database algorithms, computational methods for science and mathematics, information theory and signal processing, and more. Algorithms can either be grasped as mathematical constructs, or conceived as the implementation of such mathematical constructs into a technology, e.g. a program, or understood as an application of the technology designed for a particular task, e.g. a configuration (Mittelstadt et al. 2016). A number of such algorithms have already transformed, or reshaped, the environment of people's interaction and their daily lives. From getting insurance to landing credit, from going to college to finding a job, down to the use of search engines and GPS for navigation, interacting with the voice recognition features on our smartphones, algorithms used by computers may not only untangle very human questions. They can explain how to have better hunches and when to leave things to chance. For instance, according to the analysis of *Algorithms to Live By* (Christian and Griffiths 2016), algorithms shed light on how to deal with overwhelming choices, such as how to optimally choose a secretary, or how to optimally find parking, or a partner, by striking the right balance between stopping too early and stopping too late, namely, that which mathematicians call an "optimal stopping" problem. Quite interestingly, the approach seems to apply to angsty freshmen cri de couer as well. "When have you met enough people to know who your best match is? And what if acquiring the data costs you that very match?" (Christian and Griffiths 2016: 9-10). At the end of the day, the precise place to draw the line in all these cases, i.e. between looking and leaping, appears to settle to 37% of the pool. This means you should look at that first percentage of people you met, choosing none, but then getting ready to leap for anyone better than all of those you dated so far!

Focus of this paper is on concerns and legal challenges brought on by the use of algorithms. A particular class of algorithms that augment or replace analysis and decision-making by humans deserve here specific consideration. They either regard the sector of data analytics, namely the use of algorithms that make sense of huge streams of data (Floridi 2012; Grindrod 2014); or the discipline of machine learning, i.e. algorithms capable to define or modify decision-making rules autonomously (Van Otterlo 2013). Such algorithms often go hand-in-hand with the breathtakingly advancements in the field of artificial intelligence ("AI"), robotics, the internet of things, and more (Pagallo 2016a). This trend of delegating decisions to both automated systems and autonomous artificial agents has raised a number of critical issues. In *Weapons of Math Destruction*, for example, Cathy O'Neil warns us of how big data analytics may increase inequality and even threaten pillars of the rule of law and democracy (O'Neil 2016). Moreover, in *The Ethics of Algorithms*, Luciano Floridi and his research fellows at the Oxford Internet Institute point out six types of concern brought on by the use of algorithms (Mittelstadt et al. 2016). Whereas the first three kinds of difficulties regard epistemic concerns,[1] the normative challenges depend either on unfair outcomes ending up with discrimination, or on transformative effects that induce threats for autonomy. Such challenges motivate a final overarching concern that can be summed up in terms of traceability, which triggers issues of responsibility and the dilemmas of automation, i.e. the acceptability of replacing or augmenting human decision-making with algorithms (Pagallo and Durante 2016a).

In light of this complex scenario, the aim of the paper is to examine the current state-of-the-art in the legal domain, i.e. 'the beat of the legal drum,' in accordance with three observables of the analysis, and their variants. First, in Section 2, attention is drawn to the manifold ways in which algorithms already have impacted today's legal systems. More particularly, taking into account Balkin's work on "the laws of an algorithmic society" (Balkin 2016), focus is on obligations of transparency, matters of due process, and accountability. Section 3 complements this US-centric analysis on drawbacks and loopholes of current legal systems with a EU perspective that gives emphasis to the data protection regulation n. 679 from 2016, or

---

[1] Such concerns have to do with inconclusive evidence leading to unjustified actions, inscrutable evidence leading to opacity, and misguided evidence leading to bias.

"GDPR." Since the latter aims to discipline the entire life cycle of information that regards the production and processing of personal data through, say, Big Data sets and algorithmic techniques, this stance appears fruitful, in order to deepen our investigation on current shortcomings and gaps of the law. Section 4 intends to sum up this comparative analysis in theoretical terms. This means, on the one hand, that we should consider both the internal and external limits of the algorithmic implementations and applications in the legal domain. On the other hand, the interplay between law and technology in an "algorithmic society" should be conceived as the interaction between competing regulatory systems that need for an institutional forum of deliberation, in which both the formation and stewardship of the formal and informal rules that regulate this crucial sector of current legal systems are at stake.

The conclusion of the analysis has to do with a threefold set of problems. First, do the implementation and application of algorithms in the legal domain necessarily concern the conditions of existence and normal functioning of the rules, values, and principles that substantiate the normative context of the law? Second, would the limits of such implementations and applications hinge on the semantic irreducibility of human decisions to mathematical constructs? And third, what is the role that social acceptability and cohesion play in these cases?

The analysis of this paper starts with the different ways in which algorithms are affecting both requirements and functions of the law, i.e. what the law is supposed to be (requirements), and what it is called to do (functions). Welcome to the algorithmic society.

2. A US Approach to the Algorithmic Society

Balkin's analysis on *The Three Laws of Robotics in the Age of Big Data* appears especially fruitful in this context. The level of abstraction of this work introduces three basic observables of the analysis. First, attention is drawn to the convergence of technological developments in the fields of artificial intelligence ("AI"), robotics, big data, and more. Second, focus is on the legal principles that should govern such technological developments, namely, obligations of transparency, due process, and accountability. Third, this approach sheds light on some specific features of the US legal tradition and its constitutional values. Each of these topics is examined separately in the next paragraphs. Then, we will be ready to appreciate the EU approach to the algorithmic society below in Section 3.

2.1. Technological Convergence

As mentioned above in the introduction, algorithms can either be conceived of as mathematical constructs, or as the implementation of such mathematical constructs into a technology, or as a configuration of the technology designed for a particular task. By examining both concerns and legal challenges brought on by the use of algorithms, however, we should widen our perspective and consider their interaction with further technological developments. For instance, according to (Cath et al. 2017), the feasibility, importance, and scalability of current AI technologies go hand-in-hand with the robust and rapid progress of "four self-reinforcing trends," that concern the improvement of more sophisticated statistical and probabilistic methods, the increasing availability of large amount of data and of cheap, enormous computational power, up to the transformation of places and spaces into IT-friendly environments, e.g. smart cities and domotics. Along these lines, Balkin similarly proposes to grasp today's transformation "from the age of the Internet to the Algorithmic Society," as the convergence of social and economic decision making by algorithms, robots, and AI agents, in which "the use of robots and AI, therefore, is just a special case of the Algorithmic Society. Big Data, too, is just a feature of the Algorithmic Society. Big Data is the fuel that runs the Algorithmic Society… To vary Kant's famous statement, algorithms without data are empty; data without algorithms are blind" (Balkin 2016: 3).

A considerable amount of studies have been devoted over the past years, to stress risks, threats and challenges of this technological convergence in the fields of big data and data analytics, artificial intelligence ("AI"), or robotics. In November 2012, a non-profit US based NGO, Human Rights Watch, released a document, *Losing Humanity*, proposing the ban of "killer robots," namely "fully autonomous weapons that could select and engage targets without human intervention." In 2015, another NGO, the Future of Life Institute, released an open letter addressing the challenges presented by AI and robotics: "Its

members—and advocates, among which Bill Gates, Elon Musk, and Stephen Hawking—are concerned that as increasingly sophisticated achievements in AI accumulate - especially where they intersect with advances in autonomous robotics technology - not enough attention is being paid to safety." More recently, the White House Office of Science and Technology Policy (OSTP) conducted a series of public workshops on questions of AI and policy in 2016, culminating with a report that addresses the many ethical issues related to AI, such as fairness, accountability, and social justice, that should be tackled with increasing transparency (OSTP 2016). While the European Parliament's Committee on Legal Affairs and the UK House of Commons have released similar reports on how we should prepare for the future of AI, an Industry Connections Program within the IEEE Standards Association, namely *The Global Initiative for Ethical Considerations in the Design of Autonomous Systems*, is currently drafting another document that stresses the "ethical concerns for autonomous and intelligent systems design." Although this latter document is still in progress, attention should be drawn to the eight sectors that are under scrutiny: (i) personal data and privacy control; (ii) the law and autonomous and intelligent systems; (iii) safety and beneficence of artificial general intelligence and super-intelligence; (iv) economics of machine automation and humanitarian issues; (v) methodologies to guide ethical research, design, and manufacturing; (vi) methodologies to imbue ethics and values into AI; (vii) lethal autonomous weapons systems; and, (viii) general principles that apply to all types of autonomous and intelligent systems regardless of whether they are physical robots, e.g. care robots or driverless cars, or software AI systems.

Against this framework, we can appreciate the specificity of the legal challenges that regard the collection, processing, use, distribution and sale of data that make algorithms work. Going back to Balkin's analysis, such challenges can be summed up in accordance with the three legal principles of the Algorithmic Society that should regulate relationships of informational power and more particularly, the asymmetry of power and knowledge between the public and private governors, and those who are governed (Durante 2015). These principles have to do with fiduciary relations and public duties of algorithmic operators that should not externalize costs and harms of their operations (Balkin 2016: 12). As substantive requirements, obligations of transparency, due process, and accountability follow as a result. What kind of obligations, rights, and duties should be instantiated, clarifies the magnitude of the technological impact of today's innovation and evolution in AI and robotics, big data and data analytics, with more and more sophisticated statistical and probabilistic methods, and cheap, enormous computational power. Such a convergence is triggering an increasing number of legal hard cases that confront us with something radically new and problematic. Can settled values and principles guide the normative context of assets and interests affected by legal decisions potentially delegated to algo-rythms in such hard cases?

2.2. The Principles of the Algorithmic Society

The first law of Balkin's algorithmic society represents "algorithmic operators" as "information fiduciaries with respect to their clients and end-users" (Balkin 2016). Such operators include large online businesses like Google, Facebook and Uber. Much as occurs with doctors and lawyers, or people who manage estates or other people's property, algorithmic operators should be conceived of as fiduciaries, because "there is a significant asymmetry in knowledge and ability between fiduciary and client, and the client can't easily monitor what the fiduciary is doing on their behalf" (*op. cit.*, at 13). These fiduciaries have two central duties, namely a duty of care, and of loyalty. Although, currently in the U.S., the law does not treat online business like fiduciaries, it is a matter of fact that robots, AI agents, and algorithms are increasingly employed in the traditional operations of such fiduciaries, as doctors, lawyers, accountants, and money managers. In light of services like Airbnb, Uber, OKCupid, Match.com, or 23 and Me, "what matters in each case is that the business induce trust and collect personal information about us and use it in ways that betray our trust and create a conflict of interest" (*op. cit.*, at 15).

The second law of Balkin's algorithmic society presents algorithmic operators, both public and private, with duties toward the general public. In the case of governments and public entities, since the latter have fiduciary obligations to the people they govern, it seems fair to admit that the use of algorithms by such public bodies makes them information fiduciaries towards the population they govern. In the case of private actors, what is at stake does not only concern the use of algorithms that can harm the end-user of a service, or of a operation, i.e. that which is covered by the first principle of the algorithmic society.

Rather, focus of the second law is on matters of extra-contractual, or tortuous, liability, that is, the use of algorithms that harm many people in society as well. Years ago, Jonathan Zittrain claimed that Facebook could use its data on end-users to manipulate them in order to swing a national election (Zittrain 2014). As a result, the second law of the algorithmic society on public duties should be complemented with the third law on the central public duty of algorithmic operators "not to be algorithmic nuisances" (Balkin 2016: 17). There are five examples of how algorithmic operators can externalize the social costs of their use of algorithms onto the general public: i.e. (i) harms to reputation through classification and risk assessment methods; (ii) discrimination; (iii) normalization via the internalization of classifications and assessments of risk through algorithms; (iv) manipulation; and, (iv) lack of due process, since "the algorithm makes decisions that affect your welfare in one of the ways noted above without transparency, an ability to monitor its operations, a means of providing rebuttal or a method of holding the algorithm accountable" (Balkin 2016: 20).

The conclusion of Balkin's analysis on the notion of algorithmic nuisance is threefold. First, "the state must decide how and when to impose regulations on business that are externalizing their costs" (Balkin 2016: 21). Second, "we must be able to identify the persons or organizations who are using the algorithm that is imposing costs on the rest of society" (*ibid*). Third, "there could be cases in which algorithmic decision-making is made by anonymous or pseudo-anonymous persons or organizations, and then the law will have to require disclosure of who is behind the algorithm in order to enforce a public duty" (*ibid*). Admittedly, we can interpret this threefold conclusion with the tenets of the Dworkinian right answer-thesis, according to which a morally coherent narrative should grasp the law in such a way that, given the nature of the legal question and the story and background of the issue, scholars can attain the answer that best justifies or achieves the integrity of the law (Dworkin 1985). The notion of algorithmic nuisance—and duties and obligations that accordingly fall upon private and public operators as information fiduciaries—can be grounded on the tenets of a morally coherent theory, such as Floridi's ethics of information (Floridi 2014). By conceiving all the agents and processes of the system in terms of information, the first moral law of this perspective claims that every form of informational entropy, i.e. any kind of impoverishment of being in the "info-sphere," ought not to be caused. Also, this informational entropy ought to be prevented or removed. Balkin's idea of "algorithmical nuisance" can thus be converted into a morally coherent theory that defines the purposes that all the norms of the system are envisaged to fulfil. Algorithmic nuisance is not but a form of informational entropy that justifies new duties and obligations for the use of algorithms (Pagallo and Durante 2009).

Yet, on the one hand, how these principles of the algorithmic society—or, what is just, or what is right, according to the laws of Floridi's ethics of information—materialise in a legal system through rules and norms is a crucial issue open to diversification. As previously mentioned in the introduction, there are clear differences between, say, the US and EU legal systems, as to how we should tackle the challenges of the algorithmic society. On the other hand, both the function and interpretation of the principles in the system do not occur in a normative vacuum but are structured by a shared set of values, i.e. a general social agreement, or not. This bifurcation is critical, because it tells us something new about the process of technological convergence and innovation from a meta-regulatory standpoint. The issues brought on by the delegation of decisions to algorithms do not only depend on the degree of predictability and reliability of such mathematical constructs. Rather, several issues hinge on the degree of social agreement, or disagreement, that characterize the normative context under examination.

Let us proceed our analysis with the specificity of both the US and the EU contexts.

2.3. Legal Contexts

The legal challenges of the algorithmic society are interdisciplinary and systemic. They are interdisciplinary because we have to do with physical threat or injury, unlawful discrimination, loss of confidentiality, identity theft, financial loss, and more, spanning from constitutional law to criminal law, business and consumer law, etc. Such challenges are systemic, because e.g. Big Data techniques of the algorithmic society regard types, rather than tokens; and hence the protection of groups, rather than individuals. Contemplate such techniques as data mining or profiling, and their models for assembling groups in accordance with certain educational, occupational or professional capabilities, social practices (e.g. a religion), or social

characteristics (e.g. an ethnicity). The aim is, among others, the prediction of people's behaviour, in order to include or exclude them from a particular service, product or credit, etc. Not only can individuals be targeted as members of a group, but they can even ignore being a part of that group on the basis of a set of ontological and epistemological predicates that cluster people into multiple categories (Pasquale 2015). Although the protection of group rights has a long history in the legal tradition (Pagallo 2017a), what is new under the legal sun of the algorithmic society concerns the kind of harm, or threat, such group protection should tackle. Going back to Balkin's investigation, "algorithms (a) construct identity and reputation through (b) classification and risk assessment, creating the opportunity for (c) discrimination, normalization, and manipulation, without (d) adequate transparency, monitoring, or due process" (Balkin 2016: 20).

As mentioned above in the previous section, Balkin suggests three principles of fair governance for the algorithmic society: both private and public algorithmic operators should be deemed as information fiduciaries with new extra-contractual obligations, and with the central public duty not to externalise the social costs of their use of algorithms. Some of these remarks are popular among European scholars as well. Balkin's comments on the risks of algorithmic decision-making without transparency, or the inability to monitor its operations, up to a lack of due process in the information era, converge with most of the European debate on the necessity to rethink such provisions as Articles 6, 8 and 14 of the European Convention of Human Rights (ECHR). In light of current trends of technological convergence in AI and robotics, big data and data analytics, cloud computing and more, "it cannot be assumed that compliance with the Data Protection Legislation will automatically ensure compliance with the fundamental rights of privacy, non-discrimination and due process in the ECHR" (Hildebrandt 2013: 48). Rather, the problem concerns a new "equality of arms during the trail and immediacy of the presentation of evidence" (*ibid.*). Likewise, Balkin's concerns on a new generation of duties and responsibilities for algorithmic operators and online providers go hand-in-hand with a comparable European debate on whether clauses of immunity for internet service providers should be amended (Pagallo 2011; Floridi and Taddeo 2016).

Still, there is a major difference between the US and EU approaches. Although, in light of the examples of such services as Uber, or Airbnb, Balkin stresses that personal information can be used in ways "that betray our trust and create a conflict of interest" (Balkin 2016: 15), he does not pay special attention to the role that data protection rules play in this context. This sort of silence may appear even bizarre to EU scholars, since many of the examples made by Balkin should be understood—and are addressed in Europe—in accordance with the provisions of the data protection framework. This different approach to the legal challenges of the algorithmic society and its informational nuisance can properly be traced back to a basic distinction between the US and EU legislations in the fields of privacy and data protection (Vladeck 2015). Whereas the EU legislation, since the 1995 directive n. 46 to the current GDPR from 2016, aims to discipline the entire life cycle of information that regards the production and processing of personal data, the US approach has been so far sectorial, as shown by such acts as the *Cable Communications Policy Act* (1984), the *Electronic Communications Privacy Act* (1986), the *Health Insurance Portability and Accountability Act* (1996), the *Identity Theft and Assumption Deterrence Act* (1998), the *CAN-SPAM Act* (2003), the *Video Voyeurism Prevention Act* (2004), and so forth. Contrary to the US federal system, therefore, the EU law offers a general framework, within which many of the legal challenges examined by Balkin should be scrutinized. This level of analysis does not mean, of course, that aside from matters of privacy and data protection, the normative challenges of the algorithmic society do not raise further legal issues, such as matters of intellectual property and data ownership, business and consumer law, fair trial, etc. However, this level of analysis on the EU data protection framework seems particularly fruitful in this context. On the one hand, it sheds light on some specific provisions regulating the use of algorithms in Europe; on the other, the analysis of such regulation deepens our comprehension of the differences between US and EU laws.

3. The EU Approach to the Algorithmic Society

Issues of data protection mostly revolve under EU law around the transparency with which such data are collected, processed and used. Individuals have the right to know the purposes for which their data are processed, as well as the right to access that data and to have it rectified. In the wording of Article 8(2) of

the EU Charter of Fundamental Rights, "such data must be processed fairly… and on the basis of the consent of the person concerned or some other legitimate basis laid down by law." This type of protection through the principles of minimization and quality of data, its controllability and confidentiality, aims to constrain the flow of information, and keep firm distinctions between individuals and society, so as to protect that which the German Constitutional Court has framed in terms of "informational self-determination" since its *Volkszählungs-Urteil* ("census decision") from 15 December 1983. This general right to the *informationelle selbst-bestimmung* of individuals includes the right to determine whether personal data can be collected and, eventually, transmitted to others; the right to determine how that data may be used and processed; the right to access that data and, where necessary, to keep it up to date; up to the right to delete that data and refuse at any time to have the data processed.

The GDPR has substantially maintained the architecture of the 1995 EU directive n. 46, which was largely based on the 1980 OECD's "information-and-consent" privacy guidelines.[2] To cut to the chase, what the GDPR aims at is strengthening both individual's rights and the powers of the European authorities, while reinforcing obligations and responsibilities of data controllers, through a directly enforceable hard law-tool, such as a EU regulation (Pagallo 2017c). This threefold dimension of the GDPR can be appreciated through some of its primary rules, such as:

- Articles 21 and 22 on individual self-determination and automated decision-making, e.g. machine learning and the algorithms capable to define or modify decision-making rules autonomously;
- Article 33 on notifications of personal data breaches to the supervisory authority "competent in accordance with Article 55," along with the powers of the latter to impose administrative fines pursuant to Article 83;
- Article 17 on the right to erasure, or the right to be forgotten, and Article 20 on data portability, as a new set of duties and obligations of data controllers.

In this context, let us restrict the focus of the analysis on Art. 22 of the GDPR. Next section aims to ascertain whether an individual right to explanation of automated decision-making exists in EU law. Then, Section 3.2 considers what kind of right to explanation would be desirable.

3.1. A Right to Explanation

A considerable amount of work has revolved around whether or not a right to explanation exists pursuant the 2016 data protection regulation in Europe (Goodman and Flaxman 2016; Hildebrandt 2016; Wachter et al. 2016). Assumed as a "promising mechanism in the broader pursuit by government and industry for accountability and transparency in algorithms, artificial intelligence, robotics, and other automated systems" (Wachter et al. 2016: 3), this legal mechanism could effectively cope with some of the challenges brought on by the use of algorithms, as stressed above in Section 2, such as harms to reputation, discrimination, normalization through the internalization of classifications and assessments of risk through algorithms, manipulation, and lack of due process. However, we should preliminarily distinguish between two kinds of explanation (Wachter et al. 2016): an *ex ante* explanation concerning the general functionality of the automated decision-making system, and an *ex post* explanation which may regard either such a system functionality, or the rationale of a specific decision taken by that smart decision-making system.

In general terms, individuals have a right to explanation which derives from the notification duties of the data controllers that shall provide the data subject with all the information necessary to ensure fair and transparent processing. In the wording of the regulation, pursuant to Articles 13(2)(f) and 14(2)(g), this information regards "the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject." As such, it seems apparent that the right to information has to do here with an *ex ante* explanation that regards the system functionality. In addition, the right to explanation of Articles 13 and 14 should be complemented with the right to explanation that stems from the right to access established by Article 15. Whilst the data subject has

---

[2] See the OECD document at
https://www.oecd.org/sti/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm (last accessed 18 January 2017).

"the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed," the right to access includes information concerning the existence of automated decision-making, the logic involved, and the impact of such processing for the data subject. Therefore, whether or not such right entails information about the rationale of a specific decision taken by a smart decision-making system, is a question that has to be deepened in connection with the rules of Article 22 on "automated individual decision-making, including profiling."

According to the first paragraph of this article, "the data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her." The wording of this article is tricky, because it refers either to general clauses, such as that which "significantly" may affect the data subject, or to vague notions on what "produces legal effects," or to unclear levels of human involvement in the decision-making process, e.g. how to interpret the formula of "a decision based solely on automated processing." Would the result of an automated processing that is not actively assessed by any human, but is formally attributed to them, fall beyond the scope of Article 22(1)?

In addition, there are good reasons to suspect that the formula of this paragraph can be read either as a prohibition, or a right to object, pursuant to the requirements established by Article 22(2). As suggested by Wachter et al. (2016), in the first case data controllers need to prove that one of such requirements on, say, consent or contracts, is met, in order to legitimately engage in automated decision-making. In the case of a right to object, the data subject shall actively object an automated decision-making process, so that the latter would persist until the data subject enters an objection. In the case of prohibition for data controllers, the right to information about matters of decision-making would regard information about contractual clauses, conditions for a legitimate consent of the data subject, or that which EU or Member States law may establish in order to protect further rights and legitimate interests. In the case of a right to object, the right to information would include "the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision" (Article 22(3)). Rather than an *ex ante* explanation concerning the general functionality of the automated decision-making system, this latter interpretation of Article 22(1) would determine an *ex post* explanation which regards the rationale of a specific decision taken by that decision-making system. Moreover, this reading of Article 22(1) aligns with the wording of Recital 71, according to which individuals "should be subject to suitable safeguards," including "an explanation of the decision reached after such assessment and to challenge the decision." Although recitals have no binding power under EU law (Klimas and Vaiciukaite 2008), the EU Court of Justice now and them uses them, so as to establish the meaning of the valid law (e.g. C-533/08 from 4 May 2010, in *TNT Express vs. AXA*).

Despite this set of open issues and uncertainties, however, we should not miss three crucial aspects of this legal framework. First, the EU regulation addresses most normative challenges of the algorithmic society, e.g. normalization through the internalization of classifications and assessments of risk through algorithms, granting a general right to an *ex ante* explanation that concerns the general functionality of the automated decision-making process. There are of course the five cases in which consent is not a condition for the lawfulness of processing, pursuant to Article 6(1)(b-f) of the GDPR. Still, Articles 13, 14, and 15 of the GDPR suggest that, as a corollary of the rights to information and access, the right to explanation, although *ex ante*, will play a main role for strengthening the protection of individual consent. Whilst the latter should be requested in an intelligible and easily accessible form, suffice it to say that, for example, PayPal's terms of service used to be longer than Shakespeare's *Hamlet*, i.e. 36275 vs. 30066 words. Matters of reputation and discrimination, manipulation and digital theft, brought on by machine learning and data analytics, are thus addressed in two different ways. Whilst the US regulations are sectorial, the EU GDPR aims to address a considerable part of these issues in general terms, i.e. in accordance with the right to the informational self-determination of individuals enshrined in Article 8 of the EU Charter of fundamental rights.

Second, another difference with the US approach should be mentioned. Whilst, in this latter legal system, scholars still debate whether and to what extent new duties and responsibilities of algorithmic operators have to be introduced, an option is to conceive them as information fiduciaries: see above in Section 2.1. On the contrary, in EU law, new duties and responsibilities of algorithmic operators have been passed upon them as data controllers. By reversing the burden of proof, for example, Article 7(1) of the

GDPR states that "where processing is based on consent, the controller shall be able to demonstrate that the data subject has consented to processing of his or her personal data." Likewise, pursuant to Article 24(1), data controllers should "be able to demonstrate that processing is performed in accordance with this Regulation," by implementing appropriate technical and organisational measures. Among them, we should mention the principle of data protection by design and by default (Article 25); the duty of notification of a personal data breach to the supervisory authorities (Article 33); and so forth. On top of that, the authorities "competent in accordance with Article 55," have the power to impose severe administrative fines pursuant to Article 83.

Third, pursuant to Article 22(2-3), an individual right to an *ex post* explanation by algorithmic operators, sub specie data controllers, has to be admitted. In accordance with the second paragraph of this article, the consent of the data subject is not required for automated processing that refers to contracts, or the protection of the data subject's rights and other freedoms and legitimate interests. In accordance with the third paragraph, when automated data processing follows that which "is necessary for entering into, or performance of, a contract between the data subject and a data controller," or alternatively is based on the data subject's explicit consent, the latter has "at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision." By comparing this wording of Article 22(3) with that of the aforementioned Recital 71—which is the only time the GDPR explicitly refers to a right to obtain an explanation for specific automated decisions—some have argued that this omission from Article 22 is critical. So that, "data subjects will not be granted a legally binding *ex post* right to explanation of automated decisions on the basis of legal safeguards in Article 22 as it currently stands" (Wachter et al. 2016: 9). However, on the one hand, it seems fair to admit that one can "contest" decisions, only on the basis of the ways in which a decision is taken. On the other hand, the same authors admit that such a right of the data subject to contest decisions "could be further buttressed by drawing on the rights to fair trial and effective remedy enshrined in Articles 6 and 13 of the European Convention on Human Rights. Without an explanation of how the algorithm works, both rights are hard to enforce, because the decisions/evidence used will be impossible to contest in court" (Wachter et al. 2016: 30).

Therefore, the legal problem of interpretation does not revolve around whether a right to an *ex post* explanation exists in EU law. Rather, the issue concerns the extent of such right, e.g. whether the right to an *ex post* explanation includes the explanation of how algorithms work. This is the final topic of this section.

3.2. How Algorithms Work

The data subject's right to an *ex post* explanation has been presented as the most fruitful legal mechanism, in order to attain accountability and transparency in algorithms, artificial intelligence, robotics, and other automated systems. The implementation of the 1995 directive in member states, however, shows the ways in which this right to explanation has been limited or more precisely, a balance had to be struck between such a right and further rights and interests protected by the law. Leaving aside the material scope of the regulation and how the latter does not apply to four sets of processing of personal data, pursuant to Article 2(2)(a-d) of the GDPR, such rights include copyright regulations, the protection of trade secrets, freedom to conduct a business, and so forth (Pagallo 2011; Wachter et al. 2016). Although data controllers would have the duty to explain to data subjects how decisions are taken through the use of their algorithms, they shall not be obligated to illustrate how the software works, what are the parameters, weights, and other technical details of the algorithm, etc. How, then, should we interpret the right to an *ex post* explanation of Article 22(3) of the GDPR? How about the balance to be struck between the data subject's right and the rights of algorithmic operators?

We can envisage three different scenarios. First, we may wait for courts and authorities to make the terms of such balance more precise. Also, there are technological ways to strike such a balance through, for example, zero knowledge proof techniques (Kroll et al. 2017), differential privacy methods (Roth and Work 2014), and more. Transparency does not seem to offer here the one-size-fits-all solution for issues of legal accountability that concern people's right to explanation vis-à-vis matters of national security, public order, trade secrets, software protection, and so forth.

Second, we may suspect that several types of processing using new technologies in artificial intelligence, robotics, and other autonomous systems, will preventively be tested in accordance with Article 35 of the GDPR and a new generation of data protection impact assessments. Whilst the aim of this provision clearly has to do with the data subject's right to an *ex ante* explanation on system functionality, such tests will make the ways in which such algorithms work clearer.

Third, we should recall the Big Data features of the algorithmic society. As stressed above in Section 2.3, Big Data techniques regard types, rather than tokens; and hence groups, rather than individuals. Significantly, pursuant to Article 80(1), the GDPR establishes that "a not-for-profit body, organisation or association" have the right to lodge complaints on behalf of data subjects. There are two exemptions in the regulation that are worth mentioning. The first way to make the collection and use of Big Data compatible with the tenets of the GDPR concerns the use of pseudonymisation techniques (see Article 4(5)). The second solution has to do with the exemption of data processing for statistical purposes, pursuant to Articles 5(1)(b) and (e), 14(5)(b), 17(3)(d), 21(6), and 89 of the GDPR. As shown by Apple's incorporation of differential privacy techniques in its data collection efforts for iOS and macOS, the overall aim is to learn as much as possible about a group while learning as little as possible about any individual in it (Roth and Work 2014).

Whether or not these provisions are good enough to strike a fair balance between rights and interests protected by the EU legal framework, is of course an open issue. However, as stressed time and again, the implementation and enforcement of the GDPR rules will not occur in a normative vacuum but rather, will be structured by a shared set of values and principles, i.e. a general social agreement, or not. Some commentators have even warned that national preferences, values, and fears will fatally shape the regulatory future of data protection in Europe (Schönberger and Padova 2016: 331). It is thus necessary to widen our perspective, by taking into account the role that social acceptability and cohesion will play in these cases. The interplay between law and technology should be conceived as the interaction between competing regulatory systems that need for an institutional forum of deliberation, in which the formation and stewardship of both the formal and informal rules that govern this crucial sector of current legal systems are at stake.

4. The Governance of the Algorithmic Society

The differences between the US and EU regulations, and proposals as to how to interpret or amend them, were under scrutiny in the previous sections. In light of these differences, the aim of this section is to point out the normative challenges of the algorithmic society that such legal systems have in common, and that may explain some of these differences. Two observables of the analysis are set by this level of abstraction, and concern: (i) both the internal and external limits of algorithms that all operators have to address; (ii) the institutional forum within which some crucial uses of algorithms and their impact on society are debated. How should we address cases of general disagreement?

4.1 Algorithms, and their Limits

There are two different kinds of limit that the implementation of algorithms into a technology, and the configuration of the technology designed for a particular social task, have to take into account. The delegation of decisions to algorithms, for example, covers neither every aspect of the law, nor all legal solutions. The difficulty concerns how to weld the different levels of automation into the complexity of the law. 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: 79). In general terms, there is the technical difficulty of programming machines, or expert systems, that abide by the law through the formalization of norms, rights, or duties (Pagallo 2016b). Legal safeguards do present highly context-dependent notions that raise a number of relevant problems when reducing the informational complexity of a legal system, in which 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: 159). Hence, we can speak here about the "internal" limits of algorithms.

On the other hand, we do not have to endorse any techno-deterministic stance to accept that, in order to catch up with the pace of science and technological innovation, legal systems are not always capable to predict and anticipate every technological change. The limits of prediction and anticipation that affect the law restrain the number of legal issues that can be delegated to algorithmic procedures. A balance has thus to be struck between automated systems and the traditional tools of the law: the aim should be to determine whether a series of tasks usually carried out through traditional means, e.g. the "ought to" of the law, can be entrusted to an algorithm (Pagallo and Durante 2016b). Both the descriptive and normative aspects of this stance suggest that which we can sum up here as the "external" limits that every algorithmic operator has to address. In the basket of legal goods, we find a necessary and even inescapable mix of automation and non-automation that is not entirely new. The history of jurisprudence and the legal tradition, after all, have put forward the distinction between an automatic interpretation and application of the law (e.g. Hart 1961), and an interpretation and application of the law that 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 twofold limit of (the implementation and configuration of) algorithms that has to take into account the regulative aim of the law: (i) as a system, which competes with other regulatory systems and moreover, (ii) as an institutional sphere in which we have to strike the fair balance between automation and non-automation.

4.2 The Institutional Dimension

The normative challenges brought on by the implementation of algorithms and their configurations do not only depend on the degree of predictability and reliability of algorithms. Rather, these issues may hinge on the degree of social agreement, or disagreement, that characterize the normative context under examination. 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, and the asymmetry of power between private and public algorithmic operators, and those who are governed. In addition, technological dependence and the corresponding grade of delegation and autonomy have to be comprehended in accordance with the set of values and principles that exist in the normative context in which the consequences of tasks and decisions delegated to automated systems are evaluated. Political choices have to be taken in the fields of robotics, AI, internet of things, and generally speaking, for the operators of the algorithmic society. From a legal point of view, however, this balance between automation and non-automation does not only concern matters of interpretation and application of what lawmakers may pass. Attention should be drawn to how we interpret these matters in accordance with the norms, values and principles of the legal system. From this latter perspective, we prevent a twofold mistake. First, we should not grasp the tension between automation and non-automation as entirely provoked by the evolution of ICTs, algorithms, and a myriad of digital technologies. Since ancient Roman times, after all, lawyers have often dealt with a complex set of notions that 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." They abound in everyday life: buying a newspaper, shopping at a mall, or on the internet, enjoying a dinner at the restaurant, etc. Here, there is an overall agreement, since "the general terms seem to need no interpretation and… the recognition of instances seems unproblematic or 'automatic'" (Hart 1961: 121).

Second, we should not think about non-automatic and automatic decisions as coextensive with the difference between humans and machines. Rather, the distinction 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, lack of time, or of information. Human behaviour frequently is not guided by conscious choices and deliberation but by the need to adapt to the environment. Going back to the field of legal interpretation, what makes a legal case "plain" is thus a commonly shared (and sufficiently clear) connection between the legal output of the case and its normative context. There is a general agreement that makes the parallel between plain cases and

automation feasible, in that 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" (Hart 1961). The implementation of algorithms, as a matter of principle, can hence 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 (Pagallo 2017b).

Contrary to the plain cases, a legal hard case concerns general disagreement that may regard the meaning of the terms framing the legal question, the ways such terms are related to each other in legal reasoning, or the role of the principles that are at stake in the case. 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 formation and stewardship of the formal and informal rules that govern the algorithmic society have thus to strike the balance between delegation of decisions to algorithms and non-delegation, and to evaluate the normative context in which the consequences of such balance will occur. In Section 2, we examined some proposals on how to grasp the nature and functions of current algorithmic operators, e.g. as information fiduciaries. According to this US-centric approach to matters of accountability, they mostly regard duties of care, and of loyalty. In Section 3, focus was on the GDPR status of algorithmic operators as data controllers. There, issues of accountability principally concerned the duties of fair processing. The relevant differences between the US and EU approaches to the duties of the algorithmic operators, however, should not obfuscate what these legal systems have in common. The key question does not revolve around whether 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 decisions: as a matter of fact, a number of intricate cognitive tasks has already been delegated to algorithms, in such fields as business (e.g. algo-wars), and the military. Therefore, the question is not how far the process of automation can go but rather, whether or not cases of social and legal disagreement can be subjected to a process of automation, and to what extent. The more ours become algorithmic societies, the more their hard cases in the legal domain need an institutional space of interpretation in which norms, values, and principles are understood and evaluated. That which is under discussion cannot be entrusted to algorithms and some 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. On this basis, we can then appreciate the different ways in which legal systems are setting up new duties for operators of the algorithmic society, namely, the differences between the US and EU laws that were examined in this paper.

5. Conclusions

The paper examined concerns and legal challenges brought on by the use of algorithms. A particular class of algorithms that augment or replace analysis and decision-making by humans, i.e. data analytics and machine learning, were under scrutiny. In Section 2, taking into account Balkin's work on "the laws of an algorithmic society," attention was drawn to obligations of transparency, matters of due process, and accountability. Section 3 complemented this US-centric analysis on drawbacks and loopholes of current legal systems with norms and principles of the GDPR. The aim has been twofold. First, the intent was to shed light on some crucial differences between the US and EU law on the regulation of algorithmic operators with their obligations and duties. Whereas the US approach is sectorial, the GDPR aims to address a considerable part of these issues in general terms, e.g. in accordance with the right to the informational self-determination of individuals enshrined in Article 8 of the EU Charter of fundamental rights. Admittedly, the legal challenges of the algorithmic society comprise further legal issues, in addition to those of data protection and informational privacy, such as matters of reputation and discrimination, manipulation and lack of due process. However, by restricting the focus of the analysis on the field of privacy and data protection, this perspective has been fruitful to stress a second basic difference between the two legal systems. In the US, scholars still debate whether and to what extent new duties and responsibilities of algorithmic operators have to amend the current framework of self-regulation and light government, as shown by the White House's OSTP report, mentioned above in Section 2.1. In EU law, much of the new duties and responsibilities of algorithmic operators have been passed upon them as data controllers.

The second purpose of the paper was to stress the set of principles and values that are stake with the normative challenges of the algorithmic society. Even with such a long articulated text, as the EU's GDPR, a new set of hard cases is fated to remain open in the legal domain: the interpretation of these texts does not only hinge on the terms framing the legal question, e.g. statistical purposes of the data processing, or on how such terms are related to each other in legal reasoning: is a right to explanation valid law in the EU? Rather, such legal hard cases will increasingly have to do with the principles that are at stake in such cases. The art of balancing automation and non-automation may end up with ways of implementation and application of algorithms that do not affect the conditions of existence and normal functioning of the rules, values, and principles that substantiate the normative context of the law. Articles 22 and 35 of the GDPR illustrate how this is possible. Yet, the normative challenges brought on by the use of algorithmic implementations and applications do not only hinge on the performance of such mathematical constructs, since these challenges often have to do with the role that social acceptability and cohesion play in these cases. This perspective casts a further light on the difference between, say, the US light government policy approach to the regulation of the algorithmic operators and the EU purpose to govern the entire life cycle of personal information. The stronger the social cohesion is, 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. Since humans will bear full responsibility for the judgment of what is socially, ethically, and legally "plain" and "hard" in social affairs, it is likely that the balance between delegation of decisions to automated systems and non-delegation will represent the main topic of the legal debate over the next years.

References

Balkin, Jack M. (2016) The Three Laws of Robotics in the Age of Big Data, October, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2890965 (last accessed 4 February 2017);

Cath, Corinne, Wachter, Sandra, Mittelstadt, Brent, Taddeo, Mariarosaria and Luciano Floridi (2016) Artificial Intelligence and the 'Good Society': the US, EU, and UK Approach, December, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2906249 (last accessed 4 February 2017);

Christian, Brain and Tom Griffiths (2016) *Algorithms to Live By: The Computer Science of Human Decisions*. Holt: New York;

Durante, Massimo (2015) The Democratic Governance of Information Societies. A Critique to the Theory of Stakeholders, *Philosophy and Technology*, 28(1): 11-32;

Ronald Dworkin (1985) *A Matter of Principle*, Oxford; Oxford University Press;

Floridi, Luciano (2012) Big Data and their Epistemological Challenge, *Philosophy & Technology*, 25(4): 435-437;

Floridi, Luciano (2014) *The Ethics of Information*. Oxford: Oxford University Press;

Goodman, Bryce and Seth Flaxman (2016) EU Regulations on Algorithmic Decision-Making and a "Right to Explanation", at https://arxiv.org/abs/1606.08813 (last accessed 4 February 2017);

Grindrod, Peter (2014) *Mathematical Underpinnings of Analytics: Theory and Applications*. Oxford: Oxford University Press;

Mittelstadt, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter and Luciano Floridi (2016) The Ethics of Algorithms: Mapping the Debate, *Big Data & Society*, July-December, 1-21;

Hart, Herbert L. A. (1961) *The Concept of Law*. Oxford: Clarendon;

Hildebrandt, Mireille (2013) *Legal Protection by Design in the Smart Grid*. Report commissioned by the Smart Energy Collective (SEC). Nijmegen: Radbound University Nijmegen;

Hildebrandt, Mireille (2016) The New Imbroglio – Living with Machine Algorithms, *The Art of Ethics in the Information Society*, at https://works.bepress.com/mireille_hildebrandt/74/ (last accessed 24 February 2017);

Klimas, Tadas and Jurate Vaiciukaite (2008) The Law of Recitals in European Community Legislation, *ILSA Journal of International & Comparative Law*, 15, at https://ssrn.com/abstract=1159604 (last accessed 4 February 2017);

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;

Kroll, Joshua A., Huey, Joanna, Barocas, Solon, Felten, Edward., Reindenberg, Joel R., Robinson, David G. and Harlan Yu, Accountable Algorithms, *University of Pennsylvania Law Review*, 165 (forthcoming 2017);

Mayer Schönberger, Viktor and Yann Padova, Regime Change? Enabling Big Data through Europe's New Data Protection Regulation, *Columbia Science and Technology Law Review*, 2016, 17, 315-335;

O'Neil, Cathy (2016) *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Random House: New York;

OSTP (2016) National Science and Technology Council Networking and Information Technology, Research and Development Subcommittee, and National Science and Technology Council Networking and Information Technology, "The National Artificial Intelligence Research and Development Strategic Plan", Washington D.C.;

Pagallo, Ugo (2011) ISPs & Rowdy Web Sites Before the Law: Should We Change Today's Safe Harbour Clauses?, *Philosophy and Technology*, 24(4): 419-436;

Pagallo, Ugo (2016a) Even Angels Need the Rules: On AI, Roboethics, and the Law. In *ECAI Proceedings*, G.A. Kaminka et al. (eds.), 209-215. IOS Press, Amsterdam;

Pagallo, Ugo (2016b) Three Lessons Learned for Intelligent Transport Systems that Abide by the Law, *JusLetter IT*, 2016, at http://jusletter-it.weblaw.ch/issues/2016/24-November-2016/three-lessons-learne_9251e5d324.html (last accessed 21 January 2017);

Pagallo, Ugo (2017a) The Group, the Private, and the Individual: A New Level of Data Protection?, in Linnet Taylor, Luciano Floridi and Bart van der Sloot (eds.), *Group Privacy: New Challenges of Data Technologies*, pp. 159-173, Dordrecht: Springer;

Pagallo, Ugo (2017b) When Morals Ain't Enough: Robots, Ethics, and the Rules of the Law, *Minds and Machines*, January;

Pagallo, Ugo (2017c) The Legal Challenges of Big Data: Putting Secondary Rules First in the Field of EU Data Protection, *European Data Protection Law Review*, 3(1): 36-46;

Pagallo, Ugo and Massimo Durante (2009) Three roads to P2P systems and their impact on business ethics, *Journal of Business Ethics*, 90(4): 551-564;

Pagallo, Ugo and Massimo Durante (2016a) The Philosophy of Law in an Information Society. In *The Routledge Handbook of Philosophy of Information*, L. Floridi (ed.), 396-407. Oxon & New York: Routledge;

Pagallo, Ugo and Massimo Durante (2016b) The Pros and Cons of Legal Automation and its Governance, *European Journal of Risk Regulation*, 7(2): 323-334;

Pasquale, Frank (2015) *The Black Box Society: The Secret Algorithms that Control Money and Information*. Cambridge, MA.: Harvard University Press;

Roth, Aaron and Cynthia Work (2014) The Algorithmic Foundations of Differential Privacy, *Foundation and Trends in Theoretical Computer Science*, 9(3-4): 211-407;

Taddeo, Mariarosaria and Luciano Floridi (2016) The Debate on the Moral Responsibilities of Online Service Providers, *Science and Engineering Ethics*, 22(6): 1575-1603;

Vladeck, David (2015) Separated by Common Goals: A U.S. Perspective on Narrowing the U.S.-EU Privacy Divide, in Artemi Rallo Lombarte and Rosario García Mahamut (eds.), *Hacia un nuevo derecho europeo de protección de datos*, pp. 207-243, Tirant lo blanch, Valencia;

Wachter, Sandra, Mittelstadt, Brent and Luciano Floridi (2016) Why a Right to Explanation of Automated Decision-Making does not exist in the General Data Protection Regulation, December, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2903469 (last accessed 4 February 2017);

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. London: Hart;

Zittrain, Jonathan (2014) Facebook Could Decide an Election Without Anyone Ever Finding Out, *New Republic*, June 1.

Algo-Rhythms and the Beat of the Legal Drum

Ugo Pagallo

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

Keywords: Algorithm; Algorithmic Society; Data Protection Law; Legal Hard Case; Right to Explanation; Technological Convergence.