

The Force Awakens: Artificial Intelligence for Consumer Law

Marco Lippi

*DISMI, University of Modena and Reggio Emilia
Via Amendola 2, 42122 Reggio Emilia, Italy*

MARCO.LIPPI@UNIMORE.IT

Giuseppe Contissa

*CIRSFID, University of Bologna
Via Galliera, 3, 40121 Bologna, Italy*

GIUSEPPE.CONTISSA@UNIBO.IT

Agnieszka Jabłonowska

*University of Lodz, Faculty of Law and Administration
Doktora Stefana Kopcińskiego 8/12, 90-033 Lodz Poland*

AJABLONOWSKA@WPJA.UNI.LODZ.PL

Francesca Lagioia

*CIRSFID, University of Bologna and
Law Department, European University Institute
Villa Salviati 156, 50133, Florence, Italy*

FRANCESCA.LAGIOIA@EUI.EU

Hans-Wolfgang Micklitz

*Robert Schuman Centre for Advanced Studies
European University Institute
Villa Schifanoia, 50133 Florence Italy*

HANS.MICKLITZ@EUI.EU

Przemysław Pałka

*Yale Law School
Center for Private Law, Information Society Project
127 Wall Street, New Haven, CT 06511 USA*

PRZEMYSLAW.PALKA@YALE.EDU

Giovanni Sartor

*CIRSFID, University of Bologna and
Law Department, European University Institute
Villa Salviati 156, 50133, Florence, Italy*

GIOVANNI.SARTOR@UNIBO.IT

Paolo Torroni

*DISI, University of Bologna
Viale del Risorgimento 2, 40136 Bologna Italy*

PAOLO.TORRONI@UNIBO.IT

Abstract

Recent years have been tainted by market practices that continuously expose us, as consumers, to new risks and threats. We have become accustomed, and sometimes even resigned, to businesses monitoring our activities, examining our data, and even meddling with our choices. Artificial Intelligence (AI) is often depicted as a weapon in the hands of businesses and blamed for allowing this to happen. In this paper, we envision a paradigm shift, where AI technologies are brought to the side of consumers and their organizations, with the aim of building an efficient and effective counter-power. AI-powered tools can support a massive-scale automated analysis of textual and audiovisual data, as well as code, for the benefit of consumers and their organizations. This in turn can lead to a better oversight of business activities, help consumers exercise their rights, and enable

the civil society to mitigate information overload. We discuss the societal, political, and technological challenges that stand before that vision.

1. Introduction

Artificial Intelligence (AI) technologies have become ubiquitous in our daily experience as consumers in the digital era (Ransbotham, Gerbert, Reeves, Kiron, & Spira, 2018). AI is more and more used by businesses in order to understand and influence consumers' habits and behaviors (Helbing, 2019). Society is gaining awareness of this practice, which is increasingly looked at with concern and uneasiness, as being a source of risks for consumers, whose privacy, autonomy and wellbeing are threatened. Anticipation and control of consumers' behaviour is made possible by the huge amounts of data – concerning the activities and the characteristics of individuals – that companies are able to collect and process (Simon, 2013). In 2015, Gartner predicted that “by 2018 half of business ethics violations would occur through improper use of big data analytics” (Gartner, 2015). Even if “half” could be deemed an exaggeration, the phenomenon is too large to be ignored.

This novel scenario poses new challenges to the fields of consumer law and policy-making (BEUC, 2018). Under the umbrella of AI, technologies related also to big data (Tene & Polonetsky, 2012; Pan, 2016), robotics (Richards & Smart, 2016; Leenes et al., 2017) and algorithms (Gal & Elkin-Koren, 2017) are being analysed from the consumer protection perspective.

Over time, various regulations have been enacted to protect consumers and data subjects (Ramsay, 2012; Kaminski, 2019) and are currently being put under pressure. It is hard for consumers to resist to potential abuses and to exercise their rights. A major reason for this failure is that for consumers it is difficult to know what data practices are implemented by companies and, therefore, to pinpoint unfair or unlawful conduct. All in all, there is an urgent need to implement new countermeasures (Scherer, 2015).

A change of mindset is needed, and regulation must be complemented with consumer *empowerment*. Here is where AI could play a crucial role: that of driving technologies able to empower consumers and their organizations, by supporting consumers in safeguarding their privacy, defending their rights, protecting them from unfair practices. A real and effective *counter-power* of consumers against producers and intermediaries needs to be brought about (Lippi et al., 2019), not to build instruments that represent *alternatives* to the law, but to overcome the difficulties for consumers and regulatory agencies in *enforcing* the law.

In order to achieve this goal, several actors have to play their part and make a significant effort in the empowerment direction. First and foremost, domain experts should initiate an inter-disciplinary dialogue between lawyers and computer scientists. The importance of combining perspectives from several areas of law has already been recognized by scholars investigating these novel challenges (Kerber, 2016; Costa-Cabral & Lyskey, 2017; Helberger, Zuiderveen Borgesius, & Reyna, 2017), and this is the time to further broaden the perspective. Researchers and academics should also strive to bring cutting-edge AI technologies to the service of consumer organizations and consumer agencies,¹ as well as individual consumers. Also businesses could be incentivized to use AI not only to maximize their profits, but also to internally monitor compliance with consumer law. Such a process

1. For first signs of this process, see, e.g., (Hunt, 2018).

will require the creation of novel societally relevant technologies. In particular, we argue that consumer-empowering AI should go beyond purely data-driven approaches: it should be capable of enhancing learning from and reasoning with background legal knowledge.

The argument proceeds in four steps. In section 2 we explain how the deployment of AI by the businesses amplified certain existing risks to consumers, and created new ones. We argue that even though changes in law might be necessary, simply amending the rules will not be sufficient. What is needed is increasing the factual power of consumers: knowledge of the law, of the market, and, crucially, the ability to act upon that knowledge. This is where AI has a role to play. In section 3 we present concrete venues and dimensions where AI-driven empowerment can occur. In particular, we claim that AI can help consumers and their organisations with (1) oversight of consumer law compliance; (2) exercise of consumer rights; and (3) empowering consumers beyond their legal standing. We discuss various models of legal intervention, reasons for regulatory shortcomings, and concrete ways in which AI could help mitigate those shortcomings. In Section 4 we exemplify the types of data that can be analysed by these tools. Those include textual data (terms of service and privacy policies), audiovisual data (ads) and code (apps collecting consumer information). We then describe existing systems developed lately as early attempts to implement consumer-empowering technologies, and the broader challenges to the development of this research field. Having discussed legal, political and technological opportunities and limitations, in Section 5 we analyze the societal and economic conditions needed for the consumer empowering AI to move from the researchers' labs to the hands of actual consumers and NGOs active in the market place. Those include fostering collaboration between researchers and practitioners, creating novel schemes of funding and, broadly speaking, changing the way we, as a society, think of AI's relation to consumers. Threats notwithstanding, the potential for empowerment is there.

2. Risks for Consumers in the Big Data Era

Power imbalances between consumers and traders are arguably as old as consumer markets. Consumers have always had a lower bargaining power, they have always been prone to being manipulated into purchases, and they have always fallen short of fully understanding the boilerplate contracts they accepted. Rules on unfair commercial practices, unfair contract terms, or mandatory rights have been put in place exactly in order to counter this imbalance (Weatherill, 2001). Such corrective measures helped, though the imbalance remains (Reich, 2016; Busch, 2016; Cohen, 2019).

In the digital economy, fueled by big data and powered by machine learning algorithms, certain old risks to consumers have been amplified or transformed, and certain new ones have emerged. It has been discussed how the corporate use of AI may lead to various socially undesirable outcomes in business-to-consumer relations (Vladeck, 2015; BEUC, 2018). These include the risk of manipulation, where micro-targeted behavioral advertising may induce consumers into purchases that are against their best interest (Calo, 2014; Mik, 2016; Grafanaki, 2016), and the exclusion of particular consumer groups from access to some goods and services (Schmitz, 2014; Citron & Pasquale, 2014), also in a discriminatory manner (Zarsky, 2014; Barocas & Selbst, 2016), due to the mass-personalization of ads, prices and displayed content. Algorithmic collusion and other anti-competitive practices

can negatively affect markets, and thus opportunities for consumers (Ezrachi & Stucke, 2016; Janka & Uhsler, 2018). Even the relatively more “traditional” threats to privacy have been amplified by pervasive consumer surveillance and AI-based profiling (Hildebrandt & Gutwirth, 2008; Spencer, 2014; Pan, 2016; Zuboff, 2015).

It is thus clear that consumers are exposed to a multiplicity of threats. Some of them are new, brought about by AI, others are existing threats that became more dangerous, due to advances in AI and more in general in information technology. From a consumer policy perspective, we should consider AI as one among many technologies companies use to maximize their profit. What consumer law and policy needs to respond to is therefore not “the emergence of corporate use of AI” *per se*, but rather a broader transformation of the marketplace, of the market actors and their abilities, in which AI, alongside other socio-technological phenomena, plays a significant role.

A response to this transformation might require changes in legislation, as explored by legal researchers and policy makers: in the EU, for the time being, primarily from a private law perspective (De Franceschi, 2016; Grundmann, 2017). This, however, is not our focus in this paper. Some changes in legislation might be necessary, but we argue that just amending the law is insufficient to counter the imbalance of power. What is also needed is increasing the factual power of consumers: knowledge of the law, of the market, and, crucially, the ability to act upon that knowledge. That is why we focus on the prospect of using AI to create tools that will empower consumers and their organizations *vis a vis* the businesses.

A common critique to this view is that consumers today have a lot to gain from the new opportunities offered by businesses in the digital economy, for instance in terms of free services. We do not want to argue that there are no benefits to consumers or society at large stemming from the corporate use of AI. Effective search engines, translation tools and personalized product recommendations are clearly valuable to consumers. However, an important distinction between AI-driven products or services, and consumer-empowering AI, is given by who benefits from AI. For instance, a consumer who sees a personalized ad for a book she might enjoy, and who makes a purchase fitting her preferences, benefits from the corporate use of AI (though accepting that some of her personal data have to be processed), and so does the seller. Conversely, a consumer alerted by an AI-powered assistant that an ad is potentially unfair, benefits from empowering AI, while the seller does not. The latter is the empowerment of consumers we focus on.

3. Venues of Consumer Empowerment

Being a consumer is as much a legal status, as it is a market role (Galbraith, 1958; Roethe, 2014). From the economic viewpoint, everyone is a consumer at different points in time: when doing groceries, buying a family car or applying for a mortgage loan. The imbalance of bargaining power, experience and information between businesses and consumers is, in some sense, an inherent feature of the modern market economy. Realization of the negative effects of these asymmetries - both on the functioning of markets and on vital interests of consumers themselves (Ramsay, 2012; Rischkowsky & Döring, 2008) - has led to the enactment of consumer protection rules in all major economies around the world, including the United States and the European Union (Bourgoignie & Trubek, 2013; Reich, Micklitz,

Rott, & Tonner, 2014; Howells & Weatherill, 2005). The expansion of consumer law has gone hand in hand with the expansion of consumer markets.²

Importantly, the direct statutory interventions consumer law is usually associated with are not a suitable or sufficient remedy to *all* risks consumers face. Only some consumer problems have been addressed by legislation. Some were addressed by legislators with poor results (Loos & Luzak, 2016), which courts had to remedy (Leone, 2014), and some, arguably, were caused by poor regulation (Ben-Shahar & Schneider, 2014). Some were avoided on purpose. Overall, however, consumer law and policy point to important areas where consumers have a role to play as actors in the market game or where otherwise find themselves in need of protection.

The distinction between goals and measures employed by consumer law and policy is relevant for identifying the venues of consumer empowerment through AI. Technology alone cannot solve problems societal in nature, and at the same time simply changing the law will not automatically change societal practices. However, there appears to be a significant untapped potential in the combination of legal and technological tools in the field of consumer protection. In particular, AI can be used to create tools that would empower consumers regarding: (1) better oversight of consumer law compliance; (2) easier exercise of consumer rights; and (3) empowering consumers beyond their legal standing. In the remainder of this section we uncover the potential *venues of consumer empowerment through AI* – first against the background of existing laws and their shortcomings, and then by focusing on other (non-legal) instruments for safeguarding consumer interests.

3.1 Empowerment for More Effective Rights and Duties

In this section, we analyze statutory interventions and contractual relations as further venues for consumer empowerment, focusing on the shortcomings that motivate the need for solutions powered by AI.

3.1.1 STATUTORY INTERVENTIONS

Direct statutory interventions in consumer markets can usually be categorized as: i) duties or prohibitions imposed on traders, or ii) mandatory rights which consumers can exercise.

A prominent example of the former are the rules on unfair commercial practices in business-to-consumer relationships. Such rules in the US are set out in the Federal Trade Commission Act and in the EU they are fully and comprehensively harmonized by Directive 2005/29/EC. Both acts forbid traders, among others, from engaging in misleading advertising. Traders' freedom to conduct their business generally remains unquestioned, but boundaries to this freedom are being set to protect consumer decision-making and the functioning of the market at large (Gomez, 2006). The sanctions for violating these rules include various public interventions, most notably injunctions by the FTC in the US and fines by consumer agencies in the EU. European law also forbids companies from collecting certain types of data about consumers and it even regulates what types of clauses must not be used in standard-term contracts. Finally, different types of information duties are imposed on traders in Europe and the US. Some of these rules have been criticized as being

2. Appendix A compiles the legal acts addressing the topics discussed in Section 3.

overreaching or ineffective, yet overall they reflect the socially desirable goals related to the protection of consumers as weaker parties (Howells, 2005; Ben-Shahar & Schneider, 2014).

Consumer rights, on the other hand, are mandatory entitlements that consumers enjoy *vis-à-vis* the traders. For example, European sales law grants consumers a right to withdraw from contracts concluded at a distance (e.g., online), of which they must be informed by the trader. Consumers who shop for goods or services online are thus granted a cooling-off period of 14 days, during which they can inspect the product and change their mind. Similarly, the General Data Protection Regulation (GDPR) grants consumers rights to explanation of an automatic decision process, to object to personal data processing (including for the purposes of profiling), or to withdraw the consent which they had earlier granted. Of course, there is often a synergy between traders' duties and individual rights and consumers' initiatives to exercise their rights can positively affect traders' compliance with their corresponding obligations. The aforementioned distinction is not meant to be a clear-cut categorization, but rather to serve as lens through which one can analyze various types of regulatory shortcomings, where AI can be used to empower consumer and civil society. Importantly, many shortcomings identified with statutory interventions can equally apply to contractual settings, which are briefly discussed in the subsequent section.

3.1.2 CONTRACTUAL RELATIONS

It is worthy of note that important rights enjoyed by consumers may be granted not directly via statutes, as for the EU right to withdraw, but rather by contract. A prominent example are opt-outs of arbitration agreements, often hidden deep into the terms of service. For example, the terms and conditions set by Headspace (a popular meditation app) state (Headspace, 2019):

Opt-out. YOU MAY OPT-OUT OF THIS ARBITRATION AGREEMENT. If you do so, neither you nor Headspace can force the other to arbitrate. To opt-out, you must notify Headspace in writing no later than 30 days after first becoming subject to this arbitration agreement. Your notice must include your name and address, and the email address you used to set up your Headspace account (if you have one), and an unequivocal statement that you want to opt-out of this arbitration agreement. You must send your opt-out notice to one of the following physical or email addresses: Headspace, Inc., ATTN: Arbitration Opt-out, 2415 Michigan Avenue, Santa Monica, CA 90404; ADR@headspace.com.

In this case, it is not directly a statute, but rather the company, that gives a consumer the right to opt out. The background for this outcome remains legal: as shown in the legal practice, granting a right to opt out can be an important consideration for the courts for upholding the arbitration clause altogether.³ Still, the right remains unused if consumers are not aware of its existence (because the opt-out clause is buried under lengthy terms and conditions they do not have the skills to process), or if they are unable to exercise it for some other reasons. The clause cited above contains all information needed to exercise the right: the address, deadline for doing so, indication of what should be sent. However, once

3. See, e.g., Antonio Suarez, et al. v. Uber Technologies, Inc., No. 16-13263 (11th Cir. 2017)

again it becomes useless if consumers do not actually read the entire document, or look for such a clause.

3.1.3 CAUSES OF REGULATORY SHORTCOMINGS

One reason why duties and prohibitions imposed on traders do not achieve their set objectives is the lack of traders' compliance, which, in turn, often results from deficiencies in the enforcement system. The reasons for this state of affairs can vary. For one, factual capacities of actors involved in oversight – from state and civil society to individuals – are often limited. For example, our recent study of terms of service of online platforms indicates that all big online platforms use unfair terms of service in their consumer contracts (Lippi et al., 2019). These terms remain in use in spite of substantive regulations forbidding traders from using them. A possible explanation of this undesirable outcome is that the agencies in charge of enforcing the underlying laws simply lack the resources a thorough analysis of all terms of service would require. Only exceptionally have the relevant agencies joined forces to analyze standard terms of selected major platforms and apply collective pressure – see the example of Airbnb (European Commission, 2018).

The failure to address violations of consumer law in digital markets is also attributed to additional difficulties posed by data-driven personalization. Even the most straightforward advertising standards cannot be monitored in a systematic manner, if no-one, outside of the advertisers and their targets, really knows what ads individuals are seeing on their smartphones and computers. Therefore, unfair commercial practices become harder to detect, and appropriate countermeasures more difficult to put into action.

What also certainly impairs the effectiveness of consumer rules is the lack of consumer awareness. This becomes additionally problematic when consumers who are mostly prone to abuse - e.g. children - are those who know least about both substantive law and applicable remedies. The law might give consumers as many rights as needed, but unless the consumers know about them, and are able to make use of them, the effect of these legal entitlements will remain merely symbolic. To illustrate, under GDPR, if the processing of personal data is based on consumer's consent, such a consent should be freely given. This means that companies should not make the provision of services or the delivery of goods conditional to the consumer's consent to data processing that is unrelated to contractual performance. Consumers also enjoy the right to withdraw consent to the processing of personal data at any time and in an easy way, but are not always aware of this choice, or able to make use of it. Informing consumers about their rights (in a meaningful way) as well as providing them with easy-to-use tools for exercising their rights, is another major venue of consumer empowerment.

Finally, the non-transparency of algorithmic outcomes contributes to making the enforcement of consumers protection rules even harder. Consumers may not be able to protect themselves from unfair practices if the relevant influence remains surreptitious. For example, Facebook was infamously reported to allow targeting teenagers who felt “overwhelmed”, or “useless” - a state in which they may not critically perceive marketing communications with which they are confronted. Another prominent case concerns data processing practices, where disclosure rules are in place prompting businesses to publish privacy policies, but consumers have no way to determine whether the companies' claims match the actual

practice. The resulting inability of the consumers to detect and contest violations may provide an argument in favor of other enforcement mechanisms (e.g. via agencies). However, due to difficulties described above, these may also be unable to achieve the relevant regulatory objectives. This rightly draws attention to the possibility of addressing the “control crisis” exacerbated by algorithms by means of technological tools themselves (Bodó et al., 2017).

3.2 Empowerment beyond Legal Rights and Duties

As signalled above, consumer relations are not only domain of law, but also of economics. One of the intrinsic features of these relations is a certain imbalance of knowledge and bargaining power. Hence, the law is not the only instrument for countering this imbalance. Increases in factual capabilities of consumers (like information gathering), even if not legally mandated, are equally important tools safeguarding the proper functioning of the markets. In this section, we will show some types of empowerment that focus on the economic position of consumers, and that are not directly related to the oversight and the enforcement of legal rights and duties.

Indeed, there are many non-law-related activities that could benefit consumers, but which the latter may have limited capacity to engage in. One can think of processing data about prices, product reviews, etc. that can prove very helpful for consumers (Qazi, Syed, Raj, Cambria, Tahir, & Alghazzawi, 2016; Hong, Xu, Wang, & Fan, 2017). Even if such market-enhancing tools are not imposed by the law, consumer law and policy has long supported their development, for instance via price indication laws enhancing price comparability, or through specific tools allowing the comparison of offers in the field of electronic communications. The deployment of reliable and transparent AI systems in such domains would therefore fit the philosophy of consumer protection laws.

The discussion about empowering consumers beyond their legal standing becomes additionally relevant from a transnational perspective. Data protection offers a good illustration. The comprehensive data protection approach adopted by the EU is clearly influential beyond its borders – a phenomenon described more broadly as ‘the Brussels effect’ (Bradford, 2012) – yet a global legal standard of personal data processing has so far not been achieved. Thus the collection of particular categories of personal data may or may not be legally forbidden, depending on the applicable law. However, as the Cambridge Analytica scandal showed, even in those countries where they may be legally permissible, certain data practices may be considered socially unacceptable. Hence, merely detecting an objectionable practice and communicating about it via civil society may be of value to consumers, independently of the possibility of an effective legal recourse. Societal pressure, even in the absence of legal enforceability, might still have some influence over business behavior. This is another example of empowerment that goes beyond just relying on applicable regulations.

3.3 Conclusions on the Legal Analysis

In conclusion, even though the imbalance of power between consumers and traders is inevitable, consumer law steps in when this imbalance becomes detrimental to consumers’ interests or to the functioning of the market. More specific risks addressed by these norms include the loss of consumer autonomy, privacy violations, or undue influence – all of which

can be amplified by the corporate use of AI. Legal responses can encompass regulation of certain practices (such as data collection or advertising) and/or granting consumers actionable rights. Failure in implementing consumer protection law stems from insufficient oversight by the competent agencies, as well as from the fact that consumers may: (1) not know the law; (2) not know what is factually happening on the market; or (3) be unable to report infringements and exercise their rights. AI may benefit the implementation of consumer protection law by addressing the detection of law infringement and the assessment of compliance, as well as by supporting consumers in the exercise of their rights. For instance, AI may help consumers and their organisations in determining whether a contractual clause is unfair, whether a privacy policy violates legal requirements, whether an advertisement is potentially misleading or aggressive, etc.⁴ It can also facilitate consumers in making and aggregating their claims. Finally, AI can assist consumers in other, non-law related actions that may effectively improve their position. These include the gathering and processing of information such as product reviews and prices, the monitoring of company operations so as to exercise market pressure, and actions for claiming additional, not publicly available information from suppliers and producers regarding their products and services.

4. AI for Consumer Empowerment

Having determined *where* AI-powered tools could be used to empower consumers, we shall now elaborate on *how* AI can help consumers. In particular, we will discuss (1) the types of data AI can process to generate actionable knowledge; (2) the state of the art of consumer-empowering applications; and (3) present and future challenges.

4.1 Data Types

A major potential for AI is the ability to collect and automatically process massive amounts of data, generating knowledge that consumer organizations can use for consumer law enforcement, and consumers can rely on in order to exercise their rights. Specifically, actionable knowledge can be obtained from the (semi-)automated analysis of a variety of sources such as texts, multimedia documents, and code.

4.1.1 LARGE-SCALE TEXT ANALYTICS

Consumers are surrounded by textual documents, either produced by companies, such as terms of service, privacy policies and product specifications, or generated by users, such as product reviews. These documents contain plenty of valuable information that is especially relevant to consumers, as they may specify consumer rights and opportunities. However, it has been shown (Elshout, Elsen, Leenheer, Loos, & Luzak, 2016) that too much information can be counter-productive, to the point that consumers become unable to process the excess of information. AI can contribute to address such an “information overload” by enabling consumers to isolate and understand relevant parts and act upon them. In this way, thanks to AI, consumers could benefit from the wealth of information in their reach, just like many Internet companies do (Pałka & Lippi, 2020).

4. For a definition of ‘misleading’ and ‘aggressive’ commercial practices under EU law, see Articles 6-8 of Directive 2005/29/EC on unfair commercial practices.

4.1.2 MULTIMEDIA DATA ANALYSIS

Consumers are targeted not only with text data, but also with images, videos, games and other interactive tools. Multimedia are often used to capture the consumers' attention and to send them stimuli able to influence their behaviour, for instance through micro-targeted advertising (Zuboff, 2015; Grafanaki, 2016; Mik, 2016; Lanier, 2018).

The business model of online platforms offering services like search engines, social media sites or emails "for free" typically relies on advertising. Naturally this model creates an incentive to collect as much data as possible about consumers, for better targeting, as well as to addict consumers to the platforms, and so ensure their frequent exposure to the ads (Balkin, 2018). This leads to increasing data collection and surveillance over consumers, to the detriment of privacy, and a pervasive influence over consumers' behavior, to the detriment of autonomy (Zuboff, 2015). The decrease in consumer autonomy may lead to malfunctioning of the market, the exclusion of certain groups, etc. (Dagan, 2018). Arguably, law governing advertising needs to be updated (Ebers, 2018). However, even if we were endowed with perfect substantive rules, the problems would remain if there was no effective oversight. Unfortunately, a problem with oversight in the ad domain is the lack of reliable mechanisms for continuously collecting data about what ads are shown, to whom, and when. The challenge here is to define, circumscribe, and filter out sensitive, unnecessary and/or private data from the captured data stream (Bodó et al., 2017).

AI could help detecting aggressive advertising by analysing multimedia data. A possibility for effectively gathering and analysing such data could be to devise collective "ad-watch" campaigns consumers could be allowed to opt-in. Consumer devices could run software able to extract relevant elements from the ads a particular consumer is shown, such as images, videos, or audio signals like the voice of personal assistants making product recommendations. Such elements could then be further processed by AI techniques. Tasks such as image classification or object detection in video could be exploited in order to understand whether certain strategies are aggressive, misleading, or discriminatory. This could lead to AI-based applications endowed with normative components able to warn a consumer that a certain ad might not be legit.

4.1.3 CODE ANALYSIS

Finally, one could imagine an AI monitoring consumer data collection and explaining the practices to consumers. Just as with text and images, data about data collection could become actionable knowledge in the hands of consumers, who could either take an individual initiative or, thanks to AI, even collective action. Some research (Austin et al., 2018) suggests that the source code of apps installed on a smartphone could be reverse-engineered, which would enable uncovering with whom certain data is shared. Arguably, this scenario is the most challenging one, not only for the AI technologies involved, but also from a system's viewpoint, because data collected from a user device could be shared via other nodes of the network. For example, a service provider could collect data about consumers, and share it with its business partners, with the consumers none the wiser. Challenging as it may be, we believe that ensuring the consumers' ability to monitor the collection of their data is a task worthy of a research agenda.

4.2 Current State of the Art

Consumer-empowering AI is at its dawn. Yet, some applications already exist, that propose preliminary solutions to address some of the tasks introduced in the previous sections.

Research so far has mainly focused on textual document analysis. The aim of the Usable Privacy Policy Project⁵ is to enhance the understanding of the content of privacy policies by the end-users. Major project outcomes deal with information extraction, document classification, and question answering, exploiting state-of-the-art techniques in natural language processing. For example, the POLISIS framework (Harkous et al., 2018) uses convolutional neural networks to run an automated analysis of privacy policies, which enables answering natural language questions. On a slightly different perspective, the CLAUDETTE system (Lippi et al., 2019) automatically detects potentially unlawful clauses in online contracts using machine learning and natural language processing techniques, such as support vector machines coupled with syntactic features.

Another line of research has been dedicated to the development of legal ontologies: this is the case of PRIVONTO (Oltramari et al., 2018), which models the content of privacy policies, and PRONTO (Palmirani, Martoni, Rossi, Bartolini, & Robaldo, 2018), which represents the legal concepts within the General Data Protection Regulation.

Regarding the analysis of software compliance, the Mobile App Privacy System (MAPS) tool (Story et al., 2019) and the APPTRANS system (Austin et al., 2018) make use of machine learning techniques to check whether the implemented data practices comply with the declared policies. The key idea behind these approaches is to automatically analyze the data flow of the binary code of an app, comparing the information regarding the data actually shared with the information extracted from the privacy policies using natural language processing. Some theoretical analysis of the relation between the use of software APIs and the assessment of user privacy is also ongoing (Russell et al., 2019).

To the best of our knowledge, image and video analysis have not yet been applied to consumer-empowering AI, but the recent success of deep learning approaches in the computer vision domain (He, Zhang, Ren, & Sun, 2015) indicates that the time is ripe to pursue such a research direction too. In that respect, several works have recently focused on the contextual relevance and intrusiveness of multimedia content on the Internet (Mei & Hua, 2010), as well as on the classification of advertising images and videos (Hou et al., 2017).

4.3 Challenges Faced by AI

In order to implement consumer-empowering technologies capable of bringing about the benefits presented in Section 3, and of processing all the categories of data described in Section 4.1, AI has to face several challenges. In particular, we shall remark that state-of-the-art solutions in fields such as speech recognition, computer vision, and natural language processing largely consist of deep neural networks trained on huge data collections (LeCun, Bengio, & Hinton, 2015; He et al., 2015). These are all purely data-driven approaches and cannot fully enjoy the contributions that legal experts could give.

Although the existing state-of-the-art systems for consumer empowerment described in Section 4.2 have indeed produced promising results, in the scenarios we envisage AI will

5. <https://www.usableprivacy.org/>

need to move some steps further. In particular, we believe that a crucial contribution to the development of consumer-empowering technologies will come from the integration of deep networks with methods for knowledge representation and reasoning. Combining sub-symbolic approaches, such as neural networks, with symbolic ones, such as logic-based frameworks, is one of the major challenges for AI in the coming years (LeCun et al., 2015). In order to properly exploit AI within technologies that aim to safeguard consumers, AI needs to make use of background knowledge of the Law. To illustrate, let us consider the problem of detecting unfair clauses in contracts and policies. In this domain, classic NLP approaches such as support vector machines or recurrent neural networks achieve good performance when dealing with Terms of Service (Lippi et al., 2019) whereas more complex documents such as privacy policies need more sophisticated approaches (Contissa et al., 2018). Traditional machine learning systems are simply trained on large collections of annotated documents, where positive (potentially unfair) clauses are distinguished from the negative ones (all the remaining sentences), without exploiting any background knowledge of the legislation. Yet, an explicit representation of such knowledge would undoubtedly be a valuable resource for the task at hand, as it would provide relevant information. Besides, the reasons why a clause should be deemed fair or unfair reside in the legislation, therefore an explicit model of the legislation is required for explaining the predictions made by the classifier. A combination of symbolic and sub-symbolic methods could then increase both the effectiveness of the systems and the interpretability and explainability of deep networks, which otherwise would act as “black-box” models (Doshi-Velez et al., 2017; Miller, 2018).

Research carried out in advanced areas such as Neural-Symbolic (NeSy) learning and reasoning (Garcez et al., 2015), and Statistical Relational Learning (SRL) (Getoor & Taskar, 2007; De Raedt, Kersting, Natarajan, & Poole, 2016) has attempted diverse combinations of neural networks and statistical classifiers with logic-based formalism coming from relational learning. Frameworks such as DeepProbLog (Manhaeve, Dumancic, Kimmig, De-meester, & De Raedt, 2018), Logic Tensor Networks (Serafini & d’Avila Garcez, 2016), and Ground-Specific Markov Logic Networks (Lippi & Frasconi, 2009), to mention a few, are examples of systems that enable inferences over a set of random variables, whose outcome can be computed via a (deep) neural network, and whose inter-relations can be modeled via rules and constraints expressed in a logic formalism, which in turn can easily encode background knowledge of the domain. From a slightly different perspective, a growing number of approaches have recently been developed to exploit purely sub-symbolic (neural) approaches for reasoning tasks: this is the case, for example, of Memory Networks (Weston, Chopra, & Bordes, 2014; Sukhbaatar, Szlam, Weston, & Fergus, 2015), Neural Tensor Networks (Socher, Chen, Manning, & Ng, 2013), Neural Turing Machines (Graves, Wayne, & Danihelka, 2014) and the Neural Reasoner (Peng, Lu, Li, & Wong, 2015).

A challenge for consumer-empowering AI would then be to be able to digest and encode in such models and frameworks the large amount of information given by consumer law in the form of regulations, judgments, cases, legal ontologies. Efficiently performing reasoning over this amount of knowledge, and translating the results of this inferential process into practical recommendations and alerts for consumers would be a major achievement of this line of research.

Recently, some works have been dedicated to the classification and understanding of images and videos in ads (Vo, Tran, & Le, 2017; Hussain, Zhang, Zhang, Ye, Thomas, Agha,

Ong, & Kovashka, 2017). Here the challenge is to move from a standard categorization problem to the task of automatically detecting unfair or misleading behaviour from the provider. Clearly, this setting needs the development of novel corpora with the proper annotations. This process will have to take into account the problems of bias and fairness in machine learning datasets (Friedler et al., 2019). Since the construction of new datasets is a costly procedure, which often requires guidelines that are hard to define, another challenge for consumer-empowering AI is to exploit large unsupervised data collections. A recent approach, for example, focuses on learning specific language models and word embeddings from a corpus of unlabeled privacy policies (Harkous et al., 2018).

5. From Academic Research to Empowerment

So far, we have discussed how the realization of consumer-empowering technologies requires an advancement in the state-of-the-art for the whole research field of AI and law. Nevertheless, endowing consumers and their organizations with easy-to-use tools also requires facing societal, political and technological challenges.

To illustrate this point, we will first offer glimpses of a possible world where consumers and their organizations are empowered, and then isolate the “AI-independent” challenges that must be addressed.

Picture an app, installed on a consumer’s phone/computer, which automatically imports and analyzes the terms of service and privacy policies said consumer accepts by visiting online websites. This analysis could be conducted either in search of infringement (Lippi et al., 2019), or to assess the terms in accordance with a non-legal standard, or to inform the consumer about the rights she is granted by these contracts. Imagine also that the consumer is given an easy way to act upon the finding. For example, when she visits a website, a notification pops up saying: “I have detected 5 potentially unfair clauses. Would you like me to send the relevant data to the supervisory authority in this jurisdiction?”; or “the terms of service you have just accepted contain an arbitration clause, but you have a right to opt-out. Would you like me to send an opt-out to the company?” Finally, the consumer could have default settings in the app set to *always* send complaints to the supervisory authority or opting-out of arbitration, if this is what she actually prefers.

Similarly, watchdogs and supervisory authorities could have crawlers traversing the web, and analyzing each and every terms of service and privacy policy used in a given jurisdiction. These would not replace human lawyers with machines, but significantly increase the effectiveness of the lawyers’ work (Surden, 2014). Humans could concentrate on the difficult and creative tasks of assessment and evaluation of the pre-scanned and pre-structured data. A further development would encompass a sort of *consumer dashboard*, i.e., a platform, run for instance by NGOs, collecting all the information required by individual or collective actions. Such a platform would possess huge amounts of data about consumers, often sensitive data, and because of that it would represent a new threat to consumer privacy. However, recent research on privacy by design (Tamò-Larrieux, 2018) suggests that protective measures can be put in place.⁶

6. Moreover, we want to believe that such a civic-minded, pro-consumer initiative would gain from making use of privacy-enhancing techniques.

If, as we argue, technology already enables such tools, why do we not see them at work yet? And what can be done in order to make such tools a reality?

A present obstacle to the take up of such a technology is the expertise required for making use of such data and computational resources. Labeled data needed to train supervised learning algorithms requires the contribution of experts, and thus it is a highly time-consuming process. In other cases, we have discussed how additional datasets could be built via information sharing between consumers and their organizations. Crowdsourcing is another option (Nowak & R uger, 2010) and it has been recently proposed in the context of privacy policies (Wilson et al., 2018), though its use for the construction of high-quality corpora is still matter of discussion (Lease, 2011).

Yet, even assuming that data were publicly available, which mostly is not the case, and that computing resources were affordable, which is not true in general, building effective AI systems would still require considerable expertise. If this is an issue for businesses, who need to remain competitive, and thus require such competence, it is even more so true for the other actors in the consumer domain: consumers, consumer organizations, and public administrations. Therefore, specific policies have to be developed, aimed to fund inter-disciplinary research in the domain of consumer-empowering AI. Such policies are required for a sustainable transfer of technology and know-how from research institutions to consumer organizations.

At present, the public discourse around AI usually focuses on strategic, competitive, ethical and even defence issues. There are several establishments such as the Future of Life Institute,⁷ the AI Governance Programme at Oxford,⁸ and more closely related to the consumer domain, the Institute for Consumer Policy (ConPolicy).⁹ However, we see a need to determine measures to incentivize partnerships between research centers and administrations, for instance new funding programmes that specifically support projects providing consumer-empowerment tools and technologies to be used not only by consumers, but also by consumer organizations and by public administrations. We also believe that active empowerment needs strategic policies in order to become a reality. For example, a very much needed political intervention could be the incentive to start-ups working in the domain of consumer-empowering AI.

A related societal challenge has to do with making consumers aware of the problems and the tools to address them. Moreover, we cannot ignore a huge economic challenge, if we consider the trade-off between the regulations that have to be in place (one instance of that is the recent enactment of the GDPR) and the need to remain competitive in the global economy.

In some application domains other than law there have already been some proposals for building AI-based systems aimed to empower citizens. For example, in the context of smart cities and the Internet of Things, recent research has envisioned the development of the so-called *algorithmic guardians* (Zambonelli, Salim, Loke, De Meuter, & Kanhere, 2018), i.e., software tools completely under human control, able to protect us from “undesirable behavior of third party algorithms”. By using such systems, consumers could control the flow and processing of their personal data. Similarly, *AI guardians* have been recently proposed

7. <https://futureoflife.org>

8. <https://www.fhi.ox.ac.uk/governance-ai-program/>

9. <https://www.conpolicy.de/>

in (Etzioni & Etzioni, 2016) as “AI programs that examine AI programs”, hence heading to the development of AI systems that detect and study unfair use of AI. Moreover, there is a rich debate around the need to build “ethically-bounded” AI (Rossi & Mattei, 2019). The reason why the legal domain still lags behind may be the knowledge barrier. Legal knowledge is difficult, if not impossible, to access for individual consumers, and sometimes also for their organizations. AI could indeed fill this gap, by unlocking novel services and enabling novel technologies to the wide public.

6. Conclusions

The legal discourse on AI as a social and technological phenomenon has significantly changed over last decades. With the growing number of AI-powered applications used in consumer markets and beyond, AI started to be considered as a source of challenges, a phenomenon needing the attention of lawmakers and regulators. Notwithstanding the importance of these endeavours, we argue that legal scholars should not restrict themselves to combating the challenges posed by AI, but they should also embrace AI as an empowering tool for the civil society and the consumers. A few, peripheral, initiatives in this strand are already in place, but there is still a significant margin for progress.

We argued that applied AI does not have to remain in the hands of big businesses and corporations alone, but it can (and should) be brought to individual consumers, their organizations (civil society), as well as consumer agencies (public authorities), created to represent and defend consumer interests. Clearly, not all automation happening in these different categories and stages of empowerment would rely on AI. Nevertheless, AI would play a very important role in each of these tasks.

Specifically, we have distinguished between the different venues of empowerment. We argue that AI could be used for better oversight of business activities, easier exercise of consumer rights, and safeguarding their other, not necessarily law-related, interests. From a practical perspective, AI-powered tools could be employed to process large amounts of information (texts, audio-visual data, algorithms) to generate actionable knowledge.

However, none of this will happen without a political effort, since one cannot reasonably expect the costs of such research to be funded by the commercial sector. Likewise, political action cannot be born in a vacuum, but it should be fostered by awareness. In similar contexts, such as the movement for supporting net neutrality, significant steps forward have been made thanks to initiatives arising from the population and the civil society. We thus envision consumer-empowering AI to act as a flywheel of change, by promoting awareness, and in a sense democratizing AI, while enabling bottom-up initiatives and incentivizing regulatory processes that would otherwise be hampered by an apparent imbalance between citizens, or even consumer organizations, and the big businesses holding data and technology.

Acknowledgments

Agnieszka Jabłowska has been supported by the Foundation for Polish Science (FNP). Paolo Torroni has been partially supported by the H2020 Project AI4EU (g.a. 825619).

Appendix A. Regulations and Directives

- For foundational sources regarding the evolution of consumer law, see: President John F. Kennedy, Special message to Congress of 15 March 1962 on protecting consumer interest; Council resolution of 14 April 1975 on a preliminary programme of the European Economic Community for a consumer protection and information policy [1975] OJ C92/1; Resolution of United Nations General Assembly of 16 April 1985 on consumer protection, A/RES/39/248.
- Federal Trade Commission Act, 15 U.S.C. §§ 41-58, as amended.
- Council Directive 93/13/EEC of 5 April 1993 on unfair terms in consumer contracts [1993] OJ L95/29.
- Directive 98/6/EC of the European Parliament and of the Council of 16 February 1998 on consumer protection in the indication of the prices of products offered to consumers [1998] OJ L80/27.
- Directive 2005/29/EC of the European Parliament and of the Council of 11 May 2005 concerning unfair business-to-consumer commercial practices in the internal market and amending Council Directive 84/450/EEC, Directives 97/7/EC, 98/27/EC and 2002/65/EC of the European Parliament and of the Council and Regulation (EC) No 2006/2004 of the European Parliament and of the Council (Unfair Commercial Practices Directive) [2005] OJ L149/22.
- Directive 2008/48/EC of the European Parliament and of the Council of 23 April 2008 on credit agreements for consumers and repealing Council Directive 87/102/EEC [2008] OJ L133/66.
- Directive 2011/83/EU of the European Parliament and of the Council of 25 October 2011 on consumer rights, amending Council Directive 93/13/EEC and Directive 1999/44/EC of the European Parliament and of the Council and repealing Council Directive 85/577/EEC and Directive 97/7/EC of the European Parliament and of the Council [2011] OJ L304/64.
- Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) [2016] OJ L119/1.
- Directive (EU) 2018/1972 of the European Parliament and of the Council of 11 December 2018 establishing the European Electronic Communications Code (Recast) [2018] OJ L321/36.

References

- Austin, L. M., Lie, D., Sun, P. Y. P., Spilette, R., Wong, M., & D'Angelo, M. (2018). Towards dynamic transparency: The AppTrans (Transparency for Android applications) project. Tech. rep. 3203601, SSRN.

- Balkin, J. M. (2018). Fixing social media's grand bargain. *Hoover Working Group on National Security, Technology, and Law, Aegis Series Paper*.
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *Cal. L. Rev.*, 104, 671.
- Ben-Shahar, O., & Schneider, C. E. (2014). *More than you wanted to know: The Failure of Mandated Disclosure*. Princeton University Press.
- BEUC (2018). Automated decision making and artificial intelligence – a consumer perspective. Tech. rep., Bureau Européen des Unions de Consommateurs.
- Bodó, B., Helberger, N., Irion, K., Zuiderveen Borgesius, F., Moller, J., van de Velde, B., Bol, N., van Es, B., & de Vreese, C. (2017). Tackling the algorithmic control crisis – the technical, legal, and ethical challenges of research into algorithmic agents. *Yale JL & Tech.*, 19, 133.
- Bourgoignie, T., & Trubek, D. M. (2013). *Consumer law, common markets and federalism in Europe and the United States*, Vol. 2. Walter de Gruyter.
- Bradford, A. (2012). The brussels effect. *Northwestern University Law Review*, 107, 1.
- Busch, C. (2016). The future of pre-contractual information duties: from behavioural insights to big data. In *Research Handbook on EU Consumer and Contract Law*. Edward Elgar Publishing.
- Calo, R. (2014). Digital market manipulation. *Geo. Wash. L. Rev.*, 82, 995.
- Citron, D. K., & Pasquale, F. (2014). The scored society: due process for automated predictions. *Wash. L. Rev.*, 89, 1.
- Cohen, J. E. (2019). *Between Truth and Power: The Legal Constructions of Informational Capitalism*. Oxford University Press, USA.
- Contissa, G., Docter, K., Lagioia, F., Lippi, M., Micklitz, H.-W., Palka, P., Sartor, G., & Torroni, P. (2018). CLAUDETTE meets GDPR. Automating the evaluation of privacy policies using artificial intelligence.. Study Report, Funded by Bureau Européen des Unions de Consommateurs (BEUC).
- Costa-Cabral, F., & Lynskey, O. (2017). Family ties: the intersection between data protection and competition in eu law. *Common Market Law Review*, 54(1), 11–50.
- Dagan, H. (2018). Why markets: Welfare, autonomy, and the just society. *Mich. L. Rev.*, 117, 1289.
- De Franceschi, A. (2016). European contract law and the digital single market. *Cambridge: Intersentia*.
- De Raedt, L., Kersting, K., Natarajan, S., & Poole, D. (2016). *Statistical Relational Artificial Intelligence: Logic, Probability, and Computation*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Doshi-Velez, F., Kortz, M., Budish, R., Bavitz, C., Gershman, S., O'Brien, D., Schieber, S., Waldo, J., Weinberger, D., & Wood, A. (2017). Accountability of AI under the law: The role of explanation. *arXiv preprint arXiv:1711.01134*.
- Ebers, M. (2018). Beeinflussung und manipulation von kunden durch behavioral microtargeting. *MultiMedia und Recht (MMR)*, Heft 7.

- Elshout, M., Elsen, M., Leenheer, J., Loos, M., & Luzak, J. (2016). Study on consumers' attitudes towards terms and conditions (T&Cs). *Final report. European Commission*.
- Etzioni, A., & Etzioni, O. (2016). Keeping ai legal. *Vand. J. Ent. & Tech. L.*, 19, 133.
- European Commission (2018) http://europa.eu/rapid/press-release_IP-18-5809_en.htm. Press release, 20 September 2018.
- Ezrachi, A., & Stucke, M. E. (2016). *Virtual competition*. Oxford University Press.
- Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., Choudhary, S., Hamilton, E. P., & Roth, D. (2019). A comparative study of fairness-enhancing interventions in machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 329–338. ACM.
- Gal, M. S., & Elkin-Koren, N. (2017). Algorithmic consumers. *Harvard Journal of Law & Technology*, 30(2), 309.
- Galbraith, J. K. (1958). *The affluent society*. Houghton Mifflin.
- Garcez, A., Besold, T. R., De Raedt, L., Földiák, P., Hitzler, P., Icard, T., Kühnberger, K.-U., Lamb, L. C., Miikkulainen, R., & Silver, D. L. (2015). Neural-symbolic learning and reasoning: contributions and challenges. In *Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches, Stanford*.
- Gartner (2015) <https://www.gartner.com/en/newsroom/press-releases/2015-10-07-gartner-says-by-2018-half-of-business-ethics-violations-will-occur-through-improper-use-of-big-data-analytics>.
- Getoor, L., & Taskar, B. (2007). *Introduction to statistical relational learning*, Vol. 1. MIT press Cambridge.
- Gomez, F. (2006). The unfair commercial practices directive: A law and economics perspective. *European review of contract law*, 2(1), 4–34.
- Grafanaki, S. (2016). Autonomy challenges in the age of big data. *Fordham Intell. Prop. Media & Ent. LJ*, 27, 803.
- Graves, A., Wayne, G., & Danihelka, I. (2014). Neural Turing machines. *arXiv preprint arXiv:1410.5401*.
- Grundmann, S. (2017). *European contract law in the digital age*. Intersentia.
- Harkous, H., Fawaz, K., Lebre, R., Schaub, F., Shin, K. G., & Aberer, K. (2018). Polisis: Automated analysis and presentation of privacy policies using deep learning. In *27th USENIX Security Symposium*, pp. 531–548.
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034.
- Headspace (2019) <https://www.headspace.com/terms-and-conditions>. Effective and last updated January 17, 2019; last access May 14, 2019.

- Helberger, N., Zuiderveen Borgesius, F. J., & Reyna, A. (2017). The perfect match? a closer look at the relationship between eu consumer law and data protection law. *Common Market Law Review*, 54(5).
- Helbing, D. (2019). Societal, economic, ethical and legal challenges of the digital revolution: from big data to deep learning, artificial intelligence, and manipulative technologies. In *Towards Digital Enlightenment*, pp. 47–72. Springer.
- Hildebrandt, M., & Gutwirth, S. (2008). *Profiling the European citizen*. Springer.
- Hong, H., Xu, D., Wang, G. A., & Fan, W. (2017). Understanding the determinants of online review helpfulness: A meta-analytic investigation. *Decision Support Systems*, 102, 1–11.
- Hou, S., Chen, L., Tao, D., Zhou, S., Liu, W., & Zheng, Y. (2017). Multi-layer multi-view topic model for classifying advertising video. *Pattern Recognition*, 68, 66–81.
- Howells, G. (2005). The potential and limits of consumer empowerment by information. *Journal of Law and Society*, 32(3), 349–370.
- Howells, G., & Weatherill, S. (2005). *Consumer protection law* (2nd edition). Routledge.
- Hunt, S. (2018) <https://competitionandmarkets.blog.gov.uk/2018/10/24/emas-new-data-unit-exciting-opportunities-for-data-scientists/>.
- Hussain, Z., Zhang, M., Zhang, X., Ye, K., Thomas, C., Agha, Z., Ong, N., & Kovashka, A. (2017). Automatic understanding of image and video advertisements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1705–1715.
- Janka, S. F., & Uhsler, S. B. (2018). Antitrust 4.0 - the rise of artificial intelligence and emerging challenges to antitrust law. *European Competition Law Review*, 39, 112.
- Kaminski, M. E. (2019). Binary governance: Lessons from the gdpr’s approach to algorithmic accountability. *Southern California Law Review*, 92(6).
- Kerber, W. (2016). Digital markets, data, and privacy: competition law, consumer law and data protection. *Journal of Intellectual Property Law & Practice*, 11(11), 856–866.
- Lanier, J. (2018). *Ten arguments for deleting your social media accounts right now*. Henry Holt.
- Lease, M. (2011). On quality control and machine learning in crowdsourcing. In *Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence*.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 531, 436–444.
- Leenes, R., Palmerini, E., Koops, B.-J., Bertolini, A., Salvini, P., & Lucivero, F. (2017). Regulatory challenges of robotics: some guidelines for addressing legal and ethical issues. *Law, Innovation and Technology*, 9(1), 1–44.
- Leone, C. (2014). Transparency revisited—on the role of information in the recent case-law of the cjeu. *European Review of Contract Law*, 10(2), 312–325.
- Lippi, M., Contissa, G., Lagioia, F., Micklitz, H.-W., Pałka, P., Sartor, G., & Torroni, P. (2019). Consumer protection requires artificial intelligence. *Nature Machine Intelligence*, 1, 168–169.

- Lippi, M., & Frasconi, P. (2009). Prediction of protein β -residue contacts by markov logic networks with grounding-specific weights. *Bioinformatics*, *25*(18), 2326–2333.
- Lippi, M., Palka, P., Contissa, G., Lagioia, F., Micklitz, H.-W., Sartor, G., & Torroni, P. (2019). Claudette: an automated detector of potentially unfair clauses in online terms of service. *Artif. Intell. Law*, *27*(2), 117–139.
- Loos, M., & Luzak, J. (2016). Wanted: A bigger stick. on unfair terms in consumer contracts with online service providers. *Journal of consumer policy*, *39*(1), 63–90.
- Manhaeve, R., Dumancic, S., Kimmig, A., Demeester, T., & De Raedt, L. (2018). Deep-ProbLog: Neural probabilistic logic programming. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., & Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*, pp. 3753–3763. Curran Associates, Inc.
- Mei, T., & Hua, X.-S. (2010). Contextual internet multimedia advertising. *Proceedings of the IEEE*, *98*(8), 1416–1433.
- Mik, E. (2016). The erosion of autonomy in online consumer transactions. *Law, Innovation and Technology*, *8*(1), 1–38.
- Miller, T. (2018). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, *267*, 1–38.
- Nowak, S., & R uger, S. (2010). How reliable are annotations via crowdsourcing: A study about inter-annotator agreement for multi-label image annotation. In *Proceedings of the International Conference on Multimedia Information Retrieval*, pp. 557–566, New York, NY, USA. ACM.
- Oltramari, A., Piraviperumal, D., Schaub, F., Wilson, S., Cherivirala, S., Norton, T. B., Russell, N. C., Story, P., Reidenberg, J., & Sadeh, N. (2018). Privonto: A semantic framework for the analysis of privacy policies. *Semantic Web*, pp. 1–19.
- Palka, P., & Lippi, M. (2020). Big data analytics, online terms of service and privacy policies. In Vogl, R. (Ed.), *Research Handbook on Big Data Law*. Edward Elgar Publishing Ltd.
- Palmirani, M., Martoni, M., Rossi, A., Bartolini, C., & Robaldo, L. (2018). Pronto: Privacy ontology for legal reasoning. In *International Conference on Electronic Government and the Information Systems Perspective*, pp. 139–152. Springer.
- Pan, S. B. (2016). Get to know me: Protecting privacy and autonomy under big data’s penetrating gaze. *Harv. JL & Tech.*, *30*, 239.
- Peng, B., Lu, Z., Li, H., & Wong, K.-F. (2015). Towards neural network-based reasoning. *arXiv preprint arXiv:1508.05508*.
- Qazi, A., Syed, K. B. S., Raj, R. G., Cambria, E., Tahir, M., & Alghazzawi, D. (2016). A concept-level approach to the analysis of online review helpfulness. *Computers in Human Behavior*, *58*, 75–81.
- Ramsay, I. (2012). *Consumer law and policy: Text and materials on regulating consumer markets*. Bloomsbury Publishing.
- Ransbotham, S., Gerbert, P., Reeves, M., Kiron, D., & Spira, M. (2018). Artificial intelligence in business gets real. *MIT Sloan Management Review*, September, 17.

- Reich, N. (2016). Vulnerable consumers in eu law. In *The Images of the Consumer in EU Law: Legislation, Free Movement and Competition Law*. Hart Publishing.
- Reich, N., Micklitz, H.-W., Rott, P., & Tonner, K. (2014). *European consumer law*. Intersentia.
- Richards, N. M., & Smart, W. D. (2016). How should the law think about robots?. In *Robot Law*, pp. 3–22. Edward Elgar Publishing.
- Rischkowsky, F., & Döring, T. (2008). Consumer policy in a market economy considerations from the perspective of the economics of information, the new institutional economics as well as behavioural economics. *Journal of Consumer Policy*, 31(3), 285–313.
- Roethe, T. (2014). *Der Verbraucher. Rechtssoziologische Betrachtungen*. Nomos.
- Rossi, F., & Mattei, N. (2019). Building ethically bounded AI. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*.
- Russell, N. C., Schaub, F., McDonald, A., & Sierra-Rocafort, W. (2019). APIs and your privacy. Available at SSRN 3328825.
- Scherer, M. U. (2015). Regulating artificial intelligence systems: Risks, challenges, competencies, and strategies. *Harvard Journal of Law & Technology*, 29, 353.
- Schmitz, A. J. (2014). Secret consumer scores and segmentations: Separating haves from have-nots. *Mich. St. L. Rev.*, 1411.
- Serafini, L., & d’Avila Garcez, A. S. (2016). Logic tensor networks: Deep learning and logical reasoning from data and knowledge. In Besold, T. R., Lamb, L. C., Serafini, L., & Tabor, W. (Eds.), *Proceedings of the 11th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy)*, Vol. 1768 of *CEUR Workshop Proceedings*.
- Simon, P. (2013). *Too big to ignore: the business case for big data*, Vol. 72. John Wiley & Sons.
- Socher, R., Chen, D., Manning, C. D., & Ng, A. (2013). Reasoning with neural tensor networks for knowledge base completion. In *Advances in neural information processing systems*, pp. 926–934.
- Spencer, S. B. (2014). Privacy and predictive analytics in e-commerce. *New Eng. L. Rev.*, 49, 629.
- Story, P., Zimmeck, S., Ravichander, A., Smullen, D., Wang, Z., Reidenberg, J., Russell, C. N., & Sadeh, N. (2019). Natural language processing for mobile app privacy compliance. In *AAAI Spring Symposium on Privacy Enhancing AI and Language Technologies (PAL 2019)*.
- Sukhbaatar, S., Szlam, A., Weston, J., & Fergus, R. (2015). End-to-end memory networks. In *Advances in neural information processing systems*, pp. 2440–2448.
- Surden, H. (2014). Machine learning and law. *Wash. L. Rev.*, 89, 87.
- Tamò-Larrieux, A. (2018). *Designing for Privacy and Its Legal Framework*, Vol. 40 of *Issues in Privacy and Data Protection*. Springer International Publishing.
- Tene, O., & Polonetsky, J. (2012). Big data for all: Privacy and user control in the age of analytics. *Nw. J. Tech. & Intell. Prop.*, 11, xxvii.

- Vladeck, D. C. (2015). Consumer protection in an era of big data analytics. *Ohio NUL Rev.*, 42, 493.
- Vo, A. T., Tran, H. S., & Le, T. H. (2017). Advertisement image classification using convolutional neural network. In *2017 9th International Conference on Knowledge and Systems Engineering (KSE)*, pp. 197–202. IEEE.
- Weatherill, S. (2001). Justifying limits to party autonomy in the internal market – ec legislation in the field of consumer protection. In *Party Autonomy and the Role of Information in the Internal Market*, pp. 173–196. de Gruyter.
- Weston, J., Chopra, S., & Bordes, A. (2014). Memory networks. *arXiv preprint arXiv:1410.3916*.
- Wilson, S., Schaub, F., Liu, F., Sathyendra, K. M., Smullen, D., Zimmeck, S., Ramanath, R., Story, P., Liu, F., Sadeh, N., et al. (2018). Analyzing privacy policies at scale: From crowdsourcing to automated annotations. *ACM Transactions on the Web (TWEB)*, 13(1), 1.
- Zambonelli, F., Salim, F., Loke, S. W., De Meuter, W., & Kanhere, S. (2018). Algorithmic governance in smart cities: The conundrum and the potential of pervasive computing solutions. *IEEE Technology and Society Magazine*, 37(2), 80–87.
- Zarsky, T. Z. (2014). Understanding discrimination in the scored society. *Wash. L. Rev.*, 89, 1375.
- Zuboff, S. (2015). Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1), 75–89.