The Empire Strikes Back. Digital Control of Unfair Terms of Online Service[1]

Hans-W. Micklitz • Przemysław Pałka • Ioannis Panagis

Abstract The authors argue that it is possible to partly automate the process of abstract control of fairness of clauses in online consumer contracts. The authors present a theoretical and empirical argument for this claim, including a brief presentation of the software they have designed. This type of automation would not replace human lawyers, but would assist them and make their work more effective and efficient. Policy makers should direct their attention to the potential of using algorithmic techniques in enforcing the law regarding unfair contractual terms, and to facilitating research on and ultimately implementing such technologies.

Keywords unfair terms, digital, enforcement, automation, machine learning

Introduction

The argument of this paper is that it is possible to raise the efficiency of the process of abstract control of fairness[2] of contractual clauses in the terms of online services[3] through the partial automation of this process (using information technology). If successfully implemented, the automation of particular tasks will reduce the time and labour necessary to exercise abstract control, allowing consumer protection organizations to analyse and assess a substantially larger amount of terms of service, ultimately leading to a higher level of consumer protection in the Digital Single Market. Therefore, it is argued that policy-makers should concentrate on facilitating the research

---

[2] In the meaning of the Council Directive 93/13/EEC on Unfair Terms in Consumer Contracts, further referred to as "UCTD" or "the Directive".

[3] These contracts occur under different names: "terms of service", "terms and conditions", "service agreements", "statements", or simply "terms". The user either needs to explicitly state that he or she agrees with them, while creating an account; or such a document would contain a clause stating that by using the service, the user accepts and agrees with the document's content. For the sake of clarity, the expressions "terms of service" or "terms of online services" are used throughout this paper.

and development of the technologies capable of automating some of the tasks which nowadays must be performed by humans.

To prove the positive claim, the authors propose a theoretical model of tasks which need to be conducted by the consumer protection organisations before actual legal proceedings commence and present the range of technological options available in the automation of these tasks. Further, the authors provide empirical evidence that one of the tasks, i.e., reading contracts in search of potentially unfair clauses, can largely be automated. In the case of some types of clauses, including choice of law, jurisdiction, and some types of liability exclusion, the actual assessment of whether a clause is *in abstracto* unfair can also be conducted by a machine[4]. To provide empirical evidence, the authors have designed software capable of performing this task and have successfully tested it on twenty contractual agreements of online services. The software, called uTerms, is available online[5] and the reader is kindly invited to test it.

This research was conducted in response to a troubling phenomenon, namely the fact that despite the existence of unfair terms legislation (see Section 1.1. below), and despite the competence of consumer protection bodies and organisations to initiate the abstract control of contractual terms drafted for public use, unfair contractual clauses are prevalent in the terms of online services. This has recently been demonstrated in the literature (Loos & Luzak, 2016), has been confirmed by the authors' own research[6], and is currently being scrutinized by the European Commission[7]. The full explanation of this undesirable state of affairs would certainly need to take into account numerous factors. However, what certainly is the case, is that there are a considerable number of online platforms operating on the market, each using a contract that can be potentially checked, and these documents are often excessively long and lack transparency (Loos & Luzak, 2016). Conducting an abstract control of each and every one of them exceeds the current

---

[4] This is possible in the case of so called "black list" clauses, which exist in some jurisdictions.

[5] The reader is invited to download the software from http://uterms.software. Gradually, the website will also be filled with the annexes of this article, including the tables with a comparison of different terms of online services, and the complete dictionary that the software uses.

[6] A full table containing a comparison of unfair terms in different online platforms will soon be available on the project's website: http://uterms.software.

[7] European Commission Press release: The European Commission and Member States consumer authorities ask social media companies to comply with EU consumer rules. Brussels, 17 March 2017. Available at: http://europa.eu/rapid/press-release_IP-17-631_en.htm?locale=en

capabilities of consumer protection organizations. As a result, contracts containing unfair terms continue to be in use (thus the objective of the Directive, set out in art. 7.1[8], has not been met). Therefore, the problem has two layers: Consumer organizations and agencies might have the competence to combat the widespread use of unfair terms in the contracts of online platforms, but lack the resources to fulfil these competences. Thus, a considerable number of contracts containing unlawful and unfair clauses remain in use, to the detriment of consumer interests.

This interdisciplinary study merges the knowledge and skills of consumer lawyers, cyber lawyers and engineers. Apart from providing an argument both for the positive and the policy claim, it aims to bridge the gap between traditional legal scholarship, cyber law theory and practice, as well as information technology and data research The purpose of this study is not only to demonstrate to lawyers and policy makers that there is huge potential for IT to be used in law enforcement, but also to expose engineers to the types of challenges lawyers face—challenges with which they could be technically assisted. The private sector already employs artificial agents extensively in the conduct of business (Chopra & White, 2011). There is no reason why the state and civil society should not do so in law enforcement[9], obviously within the boundaries of the rule of law (Brownsword, 2016). However, a certain change of mind-set is necessary. Research should involve much closer cooperation between lawyers and engineers, in order to move from an abstract idea of technology (or law) "out there", to actual knowledge and understanding of technology and legal problems here and now. Additionally, it is the authors' firm belief that scholars should, at least for now, move away from big theoretical questions such as "Can computers replace judges?" to

---

[8] "Member States shall ensure that, in the interests of consumers and of competitors, adequate and effective means exist to prevent the continued use of unfair terms in contracts concluded with consumers by sellers or suppliers."

[9] What goes without saying is that such a usage would need to occur within the limits and boundaries of the law. We do not, by any means, advocate for an Orwellian future where the state uses information technologies to control the behavior of individuals. Actually, as the recent NSA scandal has revealed, such an image might not even be the future, but to a certain extent already exists. However, just as much as the public is rightly concerned with the massive preventive surveillance of individuals, not much opposition has been voiced against automated consumer law enforcement, in the interest of individuals, and against business. This may be because it would be the right step to take or perhaps, because it has not yet taken place or perhaps most likely for both reasons.

smaller, more practical questions such as "What legal tasks, which currently need to be performed by humans, could be undertaken by the machines?" The authors claim that there are several tasks that could be undertaken by machines, and demonstrate that at least one, e.g., reading of contracts and preparing legal material for the creative work of lawyers can already be automated. This issue is addressed in more detail in Section 2.1 below.

The results presented in this paper are preliminary findings of a bigger research project. The substantiation of the policy claim, i.e., researching and implementing IT solutions in the abstract control of fairness of contractual clauses, will necessarily happen in a series of smaller steps. It is the authors' firm belief that it will require an extensive cooperation, among the Member States, across sectors (private, public, non-governmental) and across academic disciplines (consumer law, legal theory, data-science, engineering). It will also require the effort of much more than one research team. For this reason, we wanted to make the idea public, and let anyone with the will, knowledge and resources to take it up and join the common effort. Moreover, since the first version of the software has been completed, and because it can already be used by anyone who might find it useful, we see no reason to keep it hidden, and so have made it available free of charge to anyone interested.

This article is divided in three major parts. In the first part, the problem is presented in more detail in its legal and socio-economic context. The state of law regarding unfair contractual terms is outlined (mostly for the non-legal audience), the contractual practice of online service providers analysed, and the difficulties of enforcing the former on the latter are explained. The second part contains the core contributions to the body of knowledge. Firstly, the model of tasks which need to be performed in the abstract control of fairness, and which could potentially be automated, is presented. Secondly, the theoretical explanation of how to automate one of these tasks, i.e., reading the terms of service and annotating the potentially unfair clauses, is offered. The explanation tackles both the task to be computed, and the available technological options. Finally, empirical evidence for the theoretical claim is given via a brief presentation of the software that the authors have designed. The quantitative results of its testing on 20 contracts are also offered. In the third part, the authors take a glimpse at he (possibly near) future, and sketch how other tasks involved in abstract control could be automated.

1. The Problem in Context

4

In this section, we outline the current state of unfair contractual terms law, the rationale behind it and its functioning[10]. This includes an analysis of the current market practice of online service providers and platforms' owners, in the context of the emerging digital economy and the European Commission's agenda on the Digital Single Market. The section concludes with a closer look at the problem itself, both on theoretical level, and through an empirical analysis of the terms of twenty online services which are currently in use.

1.1. Unfair Contractual Terms Law

The idea behind the law regarding unfair contractual terms is that since consumers often cannot influence the contents of B2C contracts they adhere to (they do not participate in the drafting); and since traders have much better knowledge of the market, products, the law and potential disputes, as well as better bargaining power, there is a serious risk that the traders will draft these contracts having only their own interest in mind, to the detriment of consumers. Such a state of affairs, where the factually stronger party also becomes legally stronger through a contract, is undesirable for several reasons, both deontological, like fairness and justice; and utilitarian, e.g. consumers' trust in the (common) market (Nebbia, 2007, pp. 8–21). Hence, the idea was that traders should not be allowed to use such terms in their contracts, and when they do, such terms should not be binding on the consumer.

Unfair contractual term legislation was adopted in numerous European states in the second part of the 20[th] century (Micklitz, Reich, Rott, & Tonner, 2014), and finally harmonized and introduced throughout the European Union with Directive 93/13/EEC on Unfair Terms in Consumer Contracts, which is still in effect. The Directive introduced minimum harmonization, meaning that the Member States were (and are) allowed to introduce stricter measures, which offer more protection for consumers. It also leaves a significant amount of freedom to Member States concerning the ways in which the Directive's objectives should be met (see below). As a result, how the regimes are actually designed differs considerably among the Member States (Schulte-

---

[10] Since the aim of the research is to bridge the gap between legal scholarship and engineering, the lawyerly audience is kindly asked to excuse the authors for being quite basic in this part. The aim of the article is not to offer an in-depth analysis of unfair terms law, but to offer an in-depth analysis of how the enforcement of this law could be assisted by technology. Hence, the target audience of Section 1.1. is not lawyers, but engineers. The latter is asked to excuse the authors in being rather basic in their explanation of the different technological options to lawyers, in Section 2.2.

Nölke, Twigg-Flesner, & Ebers, 2008, pp. 197–261). What matters, from the perspective of this paper, is the answers to two questions: "what does it mean for a term to be 'unfair'?" and "what are the means of ensuring that unfair terms are not being used by traders throughout the Union?"

What does it mean for a term to be unfair? According to the UCTD, a term is considered unfair, if :

> "contrary to the requirement of good faith, it causes a significant imbalance in
> the parties' rights and obligations arising under the contract, to the detriment of
> the consumer" (art. 3.1.)

This rather general definition is supplemented by an Annex containing an "indicative and non-exhaustive list of the terms which may be regarded as unfair" (art. 3.3.), serving both as exemplification of and guidance in detecting what European lawmakers had in mind. In addition, by 2013 the definition had been concretized by over thirty ECJ decisions (Micklitz & Kas, 2014)—and up to fifty at the moment of this paper's submission.

What is a consequence of a term being unfair from the point of view of a consumer who has agreed to it? The Directive requires that such a term "shall not be binding on the consumer" (art. 6.1). This basically means that if a consumer ends up in court, involved in a case against the trader, he or she might claim that such a term is not binding and the court will treat it as such. Moreover, the court can, or even must, declare it as one on its own motion (Micklitz & Reich 2014). However, litigation is costly, and not in the interest of traders or consumers. For that reason, it is possible that numerous consumers, bound by the unfair terms of the contracts they have agreed to, never end up going to court. As a result, there will be no proceedings in which such unfairness could be declared. For this reason, the UCTD not only foresees the possibility of concrete control of unfair terms in particular cases, but also of abstract control of contracts prepared for public use. The UCTD requires that, through their legal systems, Member States, ensure that:

> "adequate and effective means exist to prevent continued use of unfair terms in
> contracts concluded with consumers by sellers or suppliers" (art. 7.1.)

and:

> "the means referred to in paragraph 1 shall include provisions whereby persons
> or organizations, having a legitimate interest under national law in protecting

consumers, may take action according to the national law concerned before the courts or before competent administrative bodies for a decision as to whether contractual terms drawn up for general use are unfair, so that they can apply appropriate and effective means to prevent continued use of such terms" (art. 7.2.)

The "persons and organizations" and the "appropriate and effective means" differ quite significantly among the 28 Member States. In some cases, "persons and organisations" are public authorities, in other cases they are civil society organizations, and often they are both. In some countries the decision declaring the unfairness of terms is issued by a court, while in others it is issued by an administrative body. In some Member States, declaring a term unfair has an effect only on the particular trader, in others it affects all the traders operating within the jurisdiction. Concerning the latter, some countries have created so-called "grey lists" (for which listed terms will be assumed unfair, reversing the burden of proof) and "black lists" (where a term of such kind must always be considered unfair). Sometimes there are penalties for using such terms, and other times there are not. For the purposes of this study, these technicalities are of no importance. What matters is that in each and every Member State there are bodies with the competence to control the fairness of clauses in contracts drafted for general use and to legally challenge clauses which are considered unfair.

The divergence between law in books and law in action is the following: In books, if a trader uses unfair clauses in his or her contracts directed at EU consumers, there is always a body empowered to make them stop. As shown in the literature (Loos & Luzak, 2016), and as our own research indicates, online service providers do use such terms in their contracts, but the degree of control remains limited. One of the reasons for the limited degree of control is that the number of online services is just too high, their contracts are unclear and significantly long, and as a result, the task is beyond what consumer protection organizations and bodies can handle. The capabilities of consumer protection organizations and bodies could be enhanced by automating certain tasks these bodies must undertake, and by shifting the more mundane analyses from humans to machines, which would allow lawyers to concentrate on the creative work.

1.2. The Economy Goes Digital

The process of digitalization of socio-economic life commenced two decades ago, immediately attracting the attention cyber-legal-scholars (Johnson & Post, 1996; Lessig, 1999). It has since become an independent object of (socio-) legal analysis through a series now-canon

publications (Benkler, 2006; Brownsword, 2008; Goldsmith & Wu, 2006; Lessig, 2006; Palfrey & Gasser, 2008; Zittrain, 2008). Yet only in the last few years has it grasped the attention of mainstream private law scholarship in Europe, leading to a series of conferences and follow-up publications (De Franceschi, 2016; Grundmann & Kull, 2017 (forthcoming); Schulze & Staudenmayer, 2016). The Digital Single Market Strategy presented by Junker's Commission brought the digital age to the European floor[11]. The strategy defines the DSM as:

> "one in which the free movement of goods, persons, services and capital is ensured and where individuals and businesses can seamlessly access and exercise online activities under conditions of fair competition, and a high level of *consumer and personal data protection*, irrespective of their nationality or place of residence." [emphasis added]

The emergence of the information society created new opportunities, as well as new threats to consumers. The crucial phenomenon from this paper's point of view is the proliferation of online platforms[12] and services, which can be accessed by users either via web browsers, on personal computers, or via mobile applications. The significance of entities like Google, Facebook, Twitter and YouTube is self-explanatory, both in monetary terms and in terms of the important societal role they currently play (Trottier & Fuchs, 2015). However, apart from the big fish, the digital pond is inhabited by thousands of smaller ones. Searching, socializing, buying goods, trading stock, banking, healthcare, reading news, transport, housing – in essentially all sectors of the economy, there are currently online services operating, and each and every one has their own terms of service.

Users of online services either need to explicitly accept the terms when creating an account, or the terms contain a clause stating that a user agrees to these contracts simply by using the service.

---

[11] *A Digital Single Market Strategy for Europe*, Communication From The Commission To The European Parliament, The Council, The European Economic And Social Committee And The Committee Of The Regions, Brussels, 6.5.2015, COM(2015) 192 final.

[12] "What is an online platform?" is a question yet to be answered, both in terms of what kind of entities one wants to treat as platforms, and what the legal status of those entities would be. As interesting as the debate is, it cannot be addressed in detail in this paper. For reference, see the European Commission's Communication: *Online Platforms and the Digital Single Market Opportunities and Challenges for Europe*, Brussels, 25.5.2016COM (2016) 288 final.

Many of these documents contain terms which would be considered unfair according to UCTD. Joasia Luzak and Marco Loos have demonstrated that these documents often contain terms that "are unlikely to pass the UCTD's unfairness test" (2016). Luzak and Loos have identified six classes of potentially unfair terms that occur in the terms of online services:

(1) Terms that grant the online service provider a right to unilaterally change the terms of service;
(2) Terms that grant the online service provider a right to unilaterally change the service itself;
(3) Terms that grant the online service provider a right to unilaterally terminate the contract;
(4) Terms that exclude and/or limit online service provider's liability;
(5) International jurisdiction clauses; and
(6) Choice-of-law clauses.

In addition, Luzak and Loos pointed out the general lack of transparency in these documents, possibly in violation of art. 5 of the UCTD, which requires that the terms in consumer contracts be drafted in "plain, intelligible language". This last finding also strengthens one of the premises of this paper, namely that merely finding (not yet even assessing) the clauses in these contracts is a difficult and time-consuming task. Luzak and Loos based their findings on an analysis of four documents governing the relationship between consumers and Google, Facebook, Twitter and Dropbox respectively. Our own study, extended to twenty online platforms/services, has led to the conclusion that indeed, potentially unlawful clauses prevail in the terms of service one needs to adhere to before using them.

We have studied the terms of service of three search engines (Google, Yahoo and Bing), five social networks (Facebook, Twitter, LinkedIn, Snapchat and Academia.edu), two contract platforms (eBay and Amazon), two content providers (Netflix and Spotify), three content hosts (YouTube, Vimeo, 9gag), a cloud-drive (Dropbox) and five online gaming platforms (Linden Lab (Second Life), Rovio (Angry Birds), Supercell (Clash of Clans), Blizzard (World of Warcraft) and Zynga (Facebook games, including Farmville). Many of them contain all six categories of potentially unfair clauses, and none of them contains fewer than three.

Three reservations should be made at this stage. Firstly, the assessment of whether a particular contractual clause is *actually* unfair to a degree sufficient enough to remove it through the process of abstract control, needs to be conducted on a case-by-case basis. Depending on the type of

clause, this will be more or less straightforward. Some clauses, the so called "black list clauses", will always be unfair (e.g. international jurisdiction, choice of law, blank exclusion of liability "to the fullest extent permissible by law" (Micklitz, 2010). In these cases, even the automation of the assessment is possible. In other ones, such as the right to unilaterally terminate or modify a contract, will need to be assessed in more detail by a human. These are two different tasks (see the analysis in Section 2.1.), but allowing the enforcers to concentrate on this assessment instead of on reading is already an important step forward in increasing the efficiency of their work.

Secondly, there are obviously more types of clauses in the terms of online services which are potentially unfair than the ones used as examples in this study. This is true both in the case of "established" unfair clauses (e.g. mandatory international arbitration), and new, internet-specific ones (see the paragraph below). As fascinating as this subject matter is, this is not the subject of this paper. What the authors wish to demonstrate in this paper is that, regardless of what type of clause one tries to look for, its detection can be automated by information technologies.

Finally, this article sets aside the bigger question of how well the UCTD is suited to protect consumers in the digital era. After all, it was created in the early nineties, before the emergence of the internet and the digital society. While the overall result might be rather positive (Micklitz & Reich, 2014), there are digital-economy-specific issues that might need to be addressed through the supplements in the annex. This might include, for example, the role of personal data as "payment" for service, thereby connecting consumer law (which has a market logic) with personal data protection law (still grounded in fundamental rights logic). Also, what needs to be discussed soon is the role of copyright in the relationship between consumers and service providers. Users currently license platforms to commercially use their content (Twitter has the right to publish and sell a book with someone's tweets), hence the need to distinguish between the necessity to give a license to the platform for it to operate legally from other commercial activities. Further, the clauses excluding liability for the damages that computers might undergo due to viruses or malware present at the platform (currently always excluded) need to be analysed as potentially always unfair. Finally, the questions of rules affecting human rights: freedom of speech, of assembly, of religion etc. and the use of "digital force" when blocking accounts and deleting content, should be discussed (Palka, 2017; Wendehorst, 2016). The authors leave these insights here for anyone interested to study them, but as they hope is clear by know, this type of analysis in not the subject matter of this paper. What the authors wish to show is that whatever the type of potential unfairness exists, its detection can be automated. Offering an argument for the last claim is the ambition of the following section.

2. Information Technology in the Service of Consumer Organizations

In this part, the authors provide the argument for the positive claim of the paper, namely that it is possible to automate certain tasks which currently need to be performed by humans in the process of abstract control of fairness of clauses in consumer contracts. In the first section, the abstract model of these tasks is presented, and the potential of automation of each task briefly discussed. In the second section, the theoretical argument for the claim that one of these tasks, i.e. reading the contracts and highlighting the potentially unfair clauses, is offered. Possible technological solutions that can be used in the automation of this step are discussed. The third section takes the argument from theory to practice, and provides empirical evidence that this step can indeed be automated.

2.1.Unfair Contractual Clauses, (not) Seen from the Desk of a Consumer Lawyer

Why is it necessary to understand exactly what tasks need to be undertaken by a lawyer working on a case? As has recently been argued by Henry Surden (2014):

> "There may be a limited, but not insignificant, subset of legal tasks that are capable of being partially automated using current AI techniques. (…) Not all, or even most, of the tasks routinely performed by attorneys are automatable given the current state of AI technology (…) [however] there are subsets of legal tasks that are likely automatable under the current state of the art, provided that the technologies are appropriately matched to relevant tasks, and that accuracy limitations are understood and accounted for (…) the goal of such automation is not to replace an attorney, but rather, to act as a complement, for example in filtering likely irrelevant data to help make an attorney more efficient."

Surden's analysis concentrates on the work of attorneys and the potential of artificial intelligence to automate their work, but his argument can be extended to other legal professions and other technological solutions. His findings underline the fact that while not all work undertaken by lawyers can be performed by machines, some of it can. This can increase the efficiency of the work done by lawyers. However, automation will only be possible when one understands exactly what tasks need to be automated, what skills the tasks require, and whether or not the skills can be computed. One needs to realize that the work of a consumer protection agency does not *only* boil down to legal analysis, but comprises a series of other tasks. This part of the picture would not be considered by black-letter approaches to law, studying legal texts or judicial and administrative decisions. In such an approach, the "procedure" starts when a lawsuit is to be filed. Since the authors of this article claim that the substance of law is quite optimal, and the problem lies in the

sphere of factual capabilities of the enforcers, the model below is concerned with what happens *before* legal proceedings would be initiated.

If one considers how the process of abstract control looks from the perspective of a lawyer working in a consumer protection body, i.e. the actual tasks he or she must undertake, one can distinguish at least four phases (steps). Firstly, choosing and obtaining a contract to be checked. Secondly, reading the contract and choosing the clauses to be analysed as potentially unfair. Thirdly, actually assessing the unfairness of a particular clause, and consequently, studying the reasons and providing arguments to initiate the proceedings. Fourthly, potentially communicating with the trader and/or initiating the proceedings (administrative and judicial). What should be made clear is that this is an *abstract model*, understood as a "simplified version of the part of the world being studied" (Gowers, 2002), not a list of tasks performed in exactly the same way in every situation by each and every consumer organisation. The four phases are presented in Figure 1.

Fig. 1 The tasks which need to be undertaken before the legal proceedings aiming to establish abstract unfairness of a contractual clause, from the point of view of a lawyer working in a consumer protection organization

The easiest way of establishing the resource-wise significance of (not) having to perform each task is to imagine a situation in which a lawyer does not have to do it. In the case of the first step, it would be the difference between receiving a set of contracts to be studied on one's desk (computer) and beginning the work there, and having to devote time to search for these documents online in the first place. The problem here is not only having to access the website, which is hardly time consuming, but also researching the market and realizing that there is a new/relevant online service in the first place. In the case of the fourth step, it would be the difference between having to write an email/call the trader, explain to them why they should modify their terms of service, and if they refuse, to write a lawsuit and submit it; or having the computer do both. What matters here is the scale of the online market, where each of these tasks might have to be performed tens, hundreds, or even thousands of times, if all the terms of service are to be controlled. The potential of automation of step one and step four are further discussed in Section 3.2.

The core contribution of this paper concerns the automation of the second step, namely reading contracts and highlighting the potentially unfair terms of service (preparing the material for

step three). To appreciate the significance of the automation of this step, imagine the difference between receiving a set of documents with certain clauses highlighted and the command "please read the highlighted parts and assess if we should claim they are unfair" and simply receiving the set of documents with the command "please read all these contracts and look for clauses that you think we should claim to be unfair." Each requires a different mind-set and a different set of skills. The former requires actual legal analysis, subsuming the factual material to legal norms and *creative* legal work; the latter is much more laborious, time consuming, and at the same time much more *mechanical*, even in everyday language. To anticipate the findings of Section 2.3., in the authors' case study, when the software was set to annotate as many potentially unfair clauses as possible, choosing "too much noise" over " too much silence" (see Section 2.2.), came down to a difference between having to read 109.007 words and having to read 18.513 words (17%), while studying twenty contracts. The more contracts there are to check, the more significant the savings in word count will be. Furthermore, depending on the needs and choices of a particular enforcer who might prefer "more silence", this difference could be even more significant.

Finally, there is the question of the potential automation of step three, i.e., the assessment of whether a particular clause is *actually* unfair. Whether this could be automated or not will depend on the type of clause, and mostly on whether or not the clause under scrutiny is one from the "black list" (see Section 1.2.). If yes, step three can also be automated to a large degree, in the same way as step two. Examples of such clauses include mandatory international jurisdiction, mandatory choice of law different from the consumer's, or a blank exclusion of liability "to the highest degree permissible by law"(Micklitz, 2010). If not, this might be very difficult for three major reasons. Firstly, the circumstances that need to be taken into account are often extra-contractual factual conditions, which might be hard to find and represent in a machine-readable form by the machine itself. Secondly, the general principle of "causing significant imbalance in the parties' rights and obligations" cannot be formalized as such, but only through exemplification derived from past situations. As a result, at the moment of creating of the database, some situations, even obvious to a human lawyer, can be unforeseeable. Thirdly, there is often room for giving shape to the rules of the Directive through interpretation. ECJ case law provides ample evidence. Even legal experts cannot always perfectly foresee the outcome of a case that has been filed, and when it has been decided, lawyers often disagree with the legal outcome (Micklitz & Reich, 2014). This is because the reasoning leading to such a decision is not *fully formalizeable.* If humans cannot agree on the outcome, it is hard to imagine that such a task could be automated (see Section 2.2. below).

An example here would be a mutual right by both parties to terminate the agreement at any time. Consider a clause from Google's terms of service: "You can stop using our Services at any

time, although we'll be sorry to see you go. Google may also stop providing Services to you, or add or create new limits to our Services at any time."[13] At first sight, it might not seem unfair, for both parties enjoy equal legal power *ex contracto*. However, one should consider that from Google's point of view, losing one customer is hardly relevant, while from a user's point of view, having a Google account terminated *at any time* can come with tremendous losses. The losses would include losing one's own email address, without the possibility of informing contacts about the change; losing all the contacts stored in the Google account, without the possibility of creating a copy first; losing access to digital content bought in the Google Play store or stored in the cloud of Google Drive etc. For this reason, one could argue that given the type of relationship, the consumer should have a right to be informed some time in advance, for the contract to be fair. But while this can be *argued*, one does not know if *this is the law*. This is why it is not possible to automate a decision about whether a clause is *actually* unfair in many cases. The challenge here is not the technology, but the law and legal reasoning.

However, this should not be considered a problem, for this is exactly the task that needs to be performed by human lawyers. In line with the Surden's claim, one should not expect that information technology will completely *replace* lawyers any time soon. Its task is to *assist* lawyers and make their work more *efficient*. How their work can be made more efficient via the automation of step two is the subject of the following section.

2.2. Automation of Annotation in Theory: Challenges and Solutions

When analysing how a task could be automated in theory, one should keep in mind a fundamental distinction between a computational theory: "what is the goal, what is the problem to be solved by the computer?" and algorithmic implementation: "how to achieve the goal, how should the computer solve the problem?" (Alpaydin, 2016; Marr, 2010). In short, there are several ways in which a particular problem can be solved by a computer, and the choice needs to be made based on the costs and advantages of each possibility. The purpose of this subsection is to briefly introduce the problem which needs to be solved to automate step two of the model (reading and annotating), and then address the different technologies available to those who attempt to solve it.

The goal (computational theory) is to have the computer scan the document and to highlight the clauses which deserve the attention of a human lawyer, i.e., the potentially unfair terms, while leaving all the rest without annotations. In the considered research, the highlighted (annotated) parts

---

[13] https://www.google.com/policies/terms/

should correspond to a particular class of unfairness (as it stands now, based on the article by Luzak and Loos, we have five such categories: *unilateral change*, *unilateral termination*, *liability*, *choice of law* and *jurisdiction*)[14]. The computer needs to be equipped with a set of rules[15] (instructions) on how to proceed. This problem is addressed below. The argument then proceeds to algorithmic implementation (in this case: "where should the rules come from," and "how should they be created?").

The opportunity and the challenge is the following: Computers are good at scanning texts, and looking for a particular pattern, or a string of symbols. Every reader has experienced this – think of whenever you press ctrl+f to look for a word or a phrase in a document. For example, if one looks for the word "terms" in an article, the computer will display the results immediately. The same would be the case, at least within the human understanding of time, if the document were hundreds of pages long. Computers scan text really fast. Also, computers are good at searching multiple texts simultaneously and understanding specific search patterns. This example is also familiar to a researcher using the online databases, bibliographic or not, such as a library catalogue.

What computers are not good at, is *understanding the meaning* of these strings. For example, try to remember if you have ever looked for a phrase you vaguely remember having read somewhere. You remember what it said, but not really how it was phrased. Unless you are able to type at least a part of it into the search window, the computer will not find it. Or consider another example: Imagine you have a statute in a PDF file, and wonder how many types of obligations it imposes on traders. You can easily find how many times the word "obligation" is used, but you cannot "ask" the computer how many legal obligations result from the norms contained in its provisions. Unless the computer or the document has been equipped with some sort of semantic technology (Boer, 2009; Sartor, 2011), it will not "know" that the phrase "consumers have a right to demand X from traders" creates an obligation on the side of the traders corresponding to that right. To connect this with unfair terms legislation, consider two actual wordings of an international jurisdiction clause, taken from Google's and Facebook's terms, respectively:

---

[14] Once again, this is not to say that these are all the possible classes of unfair clauses that one could find in the terms of online services. These are classes which were treated as exemplary, in order to demonstrate that, whatever the type of clause, its annotation can be performed by a machine.

[15] "Rules" in the meaning ascribed by information technology: IF input X THEN output Y, not in the legal sense. However, as the reader can certainly see, both types resemble each other in structure.

"All claims arising out of or relating to these terms or the Services will be litigated exclusively in the federal or state courts of Santa Clara County, California, USA, and you and Google consent to personal jurisdiction in those courts." [16]

And:

"You will resolve any claim, cause of action or dispute (claim) you have with us arising out of or relating to this Statement or Facebook exclusively in the U.S. District Court for the Northern District of California or a state court located in San Mateo County, and you agree to submit to the personal jurisdiction of such courts for the purpose of litigating all such claims." [17]

For a human, it is clear that both phrases above express the same type of norm: IF there is to be litigation THEN we agree it will happen in the courts located in X. The sentences "all claims arising (…) will be litigated in" and "you will resolve any claim (…) in" express the same intention, and will bring about the same normative effect, but they are phrased differently. For a computer, the sentences are just two strings of symbols which are not the same. As one can easily imagine, the same provision can be worded in numerous other ways. The computer does not *know* what the provision *means*. In order to make the associations, it has to be "taught." It needs to know the "rules"(Han, Kamber, & Pei, 2011, pp. 355–363). This brings one to the choice of algorithmic implementation.

At the most general level, one could distinguish between man-made rules ("programming") and machine-inferred rules ("machine learning"), to follow the terminology of Alpaydin (2016). In the first approach, which the authors of this article have used in their research (see Section 2.3. below), it is up to a human being to write the rules for the computer. In the second approach, which has lately received a lot of scholarly attention, the human being would "give" the machine the learning-set, and let it infer the rules itself. Both have strengths and weaknesses, and both can be more or less useful depending on the stage of the research, the resources available and the needs of a particular user. In short, the weakness of "programming" is that the human needs to create the

---

[16] https://www.google.com/policies/terms

[17] https://www.facebook.com/legal/terms

rules him or herself, which takes time and can sometimes be difficult. However, the advantage is that a human knows exactly *what* the machine is doing and *why*, and can make direct modifications. On the other hand, in the case of machine learning, the creation of rules might be easier, provided that a sufficient training dataset exists (see below), but humans will not know what the rules are and fine-tuning these rules might be a very laborious, if not an impossible, process. Both are briefly analysed below.

Machine learning is a subcategory in the field of artificial intelligence (Alpaydin, 2014; Russell & Norvig, 2014), "concerned with the construction of programs that learn from experience" (Butterfield & Ngondi, 2016). Nowadays, more and more lawyers are researching the potential of machine learning to assist legal practice (Branting, 2017; Surden, 2014) which in turn has given rise to approximately 690 legal start-ups[18]. This is understandable since this technology has recently had numerous successes and has pushed the field of artificial intelligence forward[19]. It has also been successfully implemented in several research areas, ranging from machine translation (Costa-jussà & Fonollosa, 2015) to biology and medicine (Uzuner, Zhang, & Sibanda, 2009). However, one needs to understand when and for what types of tasks it can be successfully used.

In short, there are two necessary *preconditions* that have to be met in order for machine learning to function successfully: a sufficiently large dataset for the machine to learn from; and human labour and scientific consensus to collect the training dataset. Concerning the former: The learning dataset needs to contain several hundred, preferably several thousand instances. A recent example of this type of training is the paper by Aletras, Tsarapatsanis, Preoţiuc-Pietro, & Lampos (2016), in which the authors predicted with certain accuracy whether the European Court of Human Rights would find a violation of a human right in a specific case, based on a training set of 584 judgments. Albeit, the most impressive machine learning applications are trained on millions or billions of data. A celebrated example of this category is IBM Watson which was trained on, among other data, all Wikipedia entries (see to this end Chu-Carroll, Fan, Boguraev, et al., 2012). It is not a coincidence that the success of machine learning came together with the emergence of big data – it has only been possible recently, because only recently has enough data become available to learn from, as a result of the digitalization of more and more human and social activities (Alpaydin, 2016). To bring this theory to the research task considered here, successfully implementing the

---

[18] According to http://techindex.law.stanford.edu/ at the moment of writing.

[19] For the most recent examples, see for example: *The 6 most exciting AI advances of 2016*
http://www.techrepublic.com/article/the-6-most-exciting-ai-advances-of-2016/

machine learning approach in the annotation of potentially unfair terms, would require several hundred, if not several thousand actual contracts with all the potentially unfair clauses annotated and categorized. The creation of such a dataset exceeds the possibilities of a single research team. However, this does not mean that machine learning will not ever be useful in the automation of this task (the conditions are addressed in Section 3.1. below).

Concerning the second precondition – the *consensus* – consider examples where machine learning has been successfully implemented, and where it has failed. The former include speech recognition (transforming sounds into words) and image recognition (recognizing if there is a cat in a picture). Here, one can see that humans generally agree about what word has been uttered or whether or not there is a cat in a picture. On the other hand, humans often disagree about, for example, whether some conduct is morally right[20], or whether a person is beautiful[21]. This inability among humans to come to a consensus makes it difficult to create a dataset which a machine could learn from. As Chu-Carroll et al. (2012) note in the context of Question Answering (QA): "Previous experiments have shown that QA performance may not improve or may even degrade if sources are indiscriminately added." Or, in plain terms, it is not just a question of how much data one can leverage but data quality also plays an important role.

The importance of this theoretical precondition for the research under consideration is two-fold. Firstly, the broader the consensus among scholars, the more successful machine learning will be. Hence, one can see how it will be much easier to automate step two (annotation), since scholars would generally agree that a particular clause is about e.g. jurisdiction or termination rights; than step three (for scholars might disagree about whether a particular clause should be considered *actually* unfair). Consider, for example, the construction of the training dataset for predicting the result of a judgment (Aletras et al., 2016; Katz, Bommarito II, & Blackman, 2014). If outcome is assigned to one of two possibilities, namely in favour of or against the plaintiff, then this is a less complicated case, and one much easier to code and to train on. If the outcome has extra possible

---

[20] Consider the "trolley problem" in the context of self-driving cars, i.e., if a car needs to sacrifice someone's life to save another, what should it choose? For a good illustration see: http://moralmachine.mit.edu

[21] Consider the infamous example of a robot-judged beauty contest, in which the machine chose white people to be "more beautiful" than those with a darker skin color: https://www.theguardian.com/technology/2016/sep/08/artificial-intelligence-beauty-contest-doesnt-like-black-people

values, like for example, the court rejects the case, or awards a claim and so on, then it increases the possibility of discrepancies in opinion among the people creating the dataset. This problem becomes even more acute in legal "hard cases" or, for example, were a dataset supposed to codify how liberally the court interpreted the law. Such potential discrepancies and complexities in creating datasets can potentially influence the accuracy of machine learning.

Secondly, if the development of the technology is to serve the actual practical needs of the enforcers, and not simply to address the curiosity of researchers, these needs must be taken into account at a stage of the dataset creation. To understand this, consider the choice between "noise" and "silence" or, in technical jargon, the amount of precision and recall[22]. In short, the question which must be addressed by a particular consumer organisation that will use the technology is: "Given that 100% accuracy is impossible, does one prefer to have some false positives in order to ensure that there are no false negatives, or does one prefer no false positives, at the cost of possible false negatives?" In plain English: Does one accept having some random clauses annotated in order to be sure that all the potentially unfair clauses are annotated (high recall), or does one prefer to have only the relevant clauses annotated, accepting that some of the relevant clauses might be missed (high precision)? Clearly, the answer to the question depends on the needs of a particular enforcer. But if the results obtained from machine learning do not meet these needs, fine-tuning of machine-based rules might be very difficult.

This brings the argument to the second possible algorithmic implementation, i.e., man-made rules (manually created rules) which was the approach used by the authors of this paper. In short, in this approach, instead of feeding the computer with a data set, it is up to a human to create the rules that the machine will later follow. In other words, a human must find the recurring, relevant structures that occur in the text when a particular type of norm is expressed, and then generalize these structures in a machine-readable form. Then, the human must create a "dictionary" for the machine to use.

For example, consider a clause granting the provider a right to unilaterally terminate the service, taken from Google's and Dropbox's terms of use, respectively: "Google may also stop providing Services to you, or add or create new limits to our Services at any time" and "We also reserve the right to suspend or end the Services at any time at our discretion and without notice". In the first example, the structure that matters is "X may (…) stop providing services to you (…) at any time"; in the latter: "We (…) reserve a right to suspend or end the Services at any time/at our

---

[22] https://nlp.stanford.edu/IR-book/html/htmledition/evaluation-of-unranked-retrieval-sets-1.html

discretion/without notice". The computer should know that whenever it encounters such a structure, it should mark the sentence counting it as a potentially unfair unilateral termination clause.

What such a dictionary actually looks like, is of secondary importance for the purposes of this paper, but just give the reader a glimpse:

([(claims|claim|dispute|disputes)] []* [(litigated|settled|adjudicated)] []* [(court|courts)])	JURISDICTION
([you] [(consent|agree|accept)] []* [jurisdiction] []* [(court|courts)])	JURISDICTION
([(laws|law)] []* [(govern|apply)] []* [(dispute|disputes|claim|claims)])	CHOICE OF LAW
([(these|this|the)] [(terms|agreement|contract|statement)] []* [governed] []* [(law|laws)])	CHOICE OF LAW
([hold] []* [harmless] []* [(liability|claim)])	LIABILITY
([(is|be)] []* [(responsible|liable)] []* [(claim|claims|damage|damages)])	LIABILITY
([(may|can|might)] [(modify|change|revise)] []* [(terms|contract|agreement)])	UNILATERAL CHANGE
([(reserves|reserve)] [(a|the)] [right] []* [(change|amend|modify)] []* [(terms|agreement|contract)])	UNILATERAL CHANGE
([(may|can)] [stop] [providing] []* [you])	UNILATERAL TERMINATION
([(reserve|reserves)] [(a|the)] [right] [to] [(suspend|terminate|stop|cease)] []* [(service|services|contract|relation)])	UNILATERAL TERMINATION

Two examples of each category have been presented, though the actual dictionary contains dozens of them[23]. How does one generate the database? Through a repeated cycle of: *collection*, *intuition* and *testing*. In the research presented here, we started by actually reading the terms of several online services and marking the clauses we found potentially unfair. Then we extracted the structure, which could potentially be used by some other provider. For example, if one considers the wording of Facebook's choice of law clause: "The laws of the State of California will govern this Statement", what matters is the string "laws of X will govern this statement". This is the collection phase. Then, through linguistic intuition, we can easily suppose that a similar structure will be used by some other provider, but instead of "will govern" they will use "shall govern", or instead of "govern" will use "apply to", or instead of "statement", will use "contract", or "agreement" or "terms" etc. That is the intuition phase. Finally we have tested it on more contracts – and when the software fails to capture a potentially unfair clause, we extract that wording, and enrich it through intuition, and place it in the database.

This approach, though technologically less spectacular, has certain advantages. Its shortcoming is that the creation of the dictionary is a laborious process. However, to create a dictionary that will give "good enough" results (high precision and/or recall, depending on the needs) one needs a much smaller dataset than for machine learning. The one used by the authors of

---

[23] It can be found at the project's webpage: http://uterms.software

this paper consisted of 20 contracts (109.000 words), and led to quite good results (see Section 2.3. below). Instead, using machine learning could be optimal when it comes to rule creation, but only once a much larger dataset (containing several hundred, or thousands of documents) is available. It does not yet exist, but that does not mean that it will never come into existence (see Section 3.1. below). Secondly, fine-tuning of the rules, i.e. reducing noise to increase precision, or reducing silence to increase recall, is much easier since a human can directly modify the rules.

As has been demonstrated, the computational problem in question (having the computer annotate and categorize the potentially unfair clauses) can be solved through different algorithmic implementations. Machine learning could be used in the future, but only when a big enough dataset is available. For now, manually-created rules yield more precise results, and offer certain advantages during the usage stage. The next section offers empirical evidence for this theoretical claim.

2.3. Automation of Annotation: From Theory to Practice

In order to prove the positive claim, elaborated in theory above, the authors of this paper have designed software capable of automating the annotation, and have created a dictionary to be used by the software. The software, called uTerms, can be downloaded free of charge from the project's website, and the reader is kindly invited test it[24].

From the user's perspective, the functionalities of the software are the following. Firstly, the programme imports a contract to be scanned – it can import it either from a file (like a pdf, or docx, or html), if the user has it saved on their computer, or directly from the internet. For the latter, it is sufficient for a user to copy/paste the URL from an internet browser into the software's bar. A last option is to import the contract to the programme by directly copy/pasting the contents of the contract into the programme's main window. Then, the user can click "Process" and, as output, they will receive the same document, with the potentially unfair contractual clauses highlighted. The highlighted (annotated) parts will be in different colours, each corresponding to a particular class of unfairness. The colour coding is shown on the right column of Figure 2 under the title "Annotations". The user might choose to display only a particular type of unfairness (e.g. only choice of law, or choice of law and liability, etc.), depending on their needs. In the example shown above, the software displays the entire contract. However, the software can display solely the annotated paragraphs (thereby sharply reducing the amount of text displayed). Additionally, the

---

[24] http://uterms.software

software is capable of "batch-processing" and comparing several contracts at the same time. These results are shown in the programme's output window. The user can save the results as a pdf or docx document. All this can be seen in Figure 2, below.
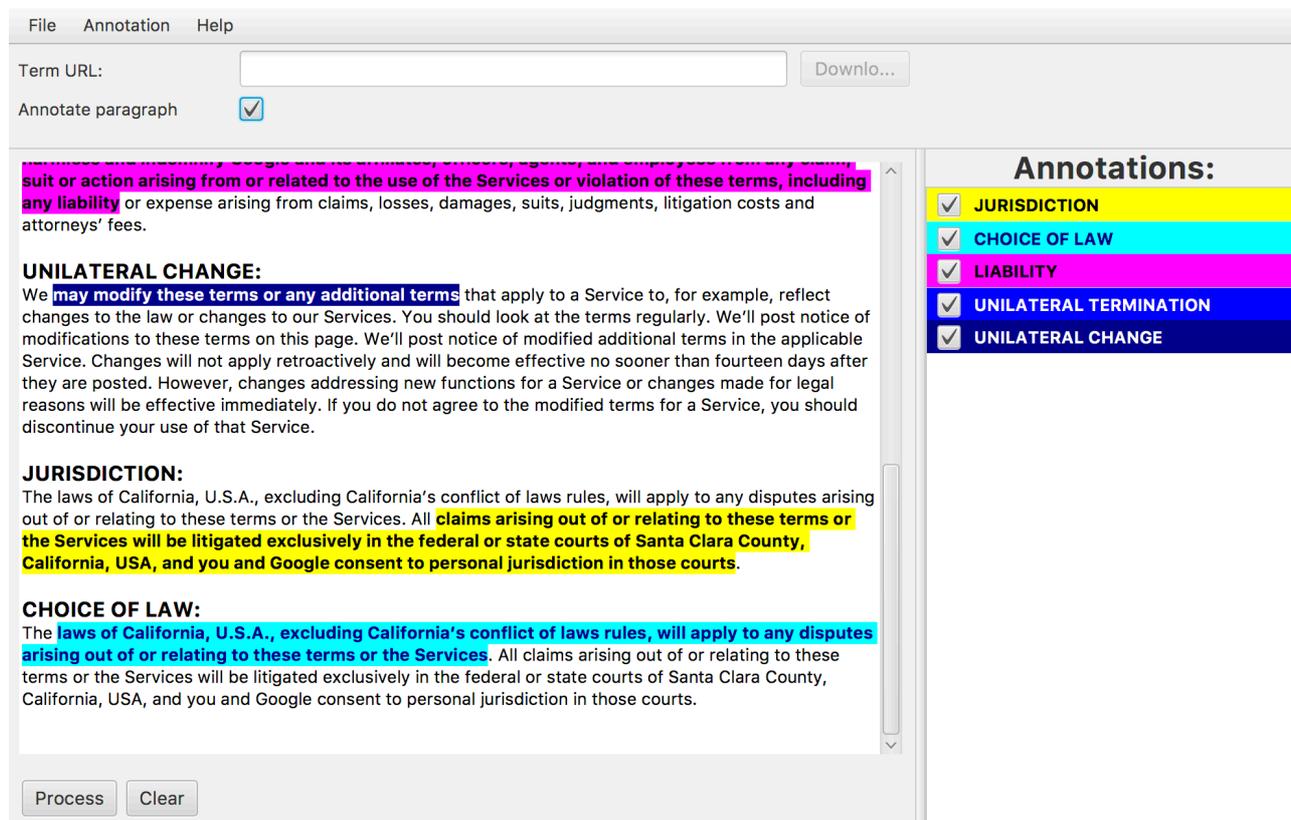

Fig. 2 The interface and the functionalities of the software

What should be made clear is that what the software will annotate depends not on its code, but on the dictionary it uses. As explained in the previous section, the dictionary itself might be fine-tuned to reflect the needs of a given enforcer. Particularly, the amount of "noise" and "silence" can be played with, depending on whether the user prefers to have higher recall or higher precision.

In the empirical part of the research, the authors chose to create a dictionary that would annotate all the potentially unfair clauses (having no false negatives), at the cost of having some noise (accepting some false positives). As a result, in a corpus of 20 contracts of the services enumerated in Section 1.2. above, the software reduced 109.007 words of unstructured text to 18.513 words (17%) of categorized and annotated text, thereby reducing the amount of text that needs to be read by 83%. A lawyer working for a consumer protection agency or association might move directly to the step three, i.e., the actual legal analysis. The number of false positives in this setting was 27 instances out of 154 annotated clauses, amounting to 18%. Hence, with recall at 1

(100%)[25] the precision is 0.82 (82%). In the authors' opinion, this is quite a successful result, given that this is only the first step, and that the corpus was not very large. However, one must remember that these numbers result not from the design of the software, but from the design choices in the construction of the dictionary, and can differ significantly, depending on the different choices one might make. Our claim is not that the functioning is already ideal – on the contrary, it can certainly be improved in many respects. This is one of the reasons we decided to make the software public now. Our claim is much more modest, but at the same time consequential: It is possible to partly automate the annotation of potentially unfair clauses in the terms of online platforms. We have shown this theoretically, and have proven it empirically. Hence the policy claim that there should be more research and actual attempts to implement information technology in the abstract control of contractual terms.

What else can be done? What are the next steps to take? Where could research interests be directed? The following section seeks to answer these questions.

3. A Peek Into the (Possibly Near) Future: What Else Does the Technology Have to Offer?

The results presented in this paper are only the first step in a bigger project, which will be continued by the authors, but might be also taken up by other researchers. The entire system of abstract control can be revolutionised with the usage of information technology, but this revolution, if it is to happen and not only be talked about, will need to take place through numerous small steps and with the collaboration of researchers from different fields, policy makers and consumer protection organisations. Here, the authors would like to signal how other steps in the process of abstract control fairness of contractual clauses might be automated (document collection and communication/initiation of proceedings, Section 3.2) and how other technologies, i.e. machine learning, can be used in the automation of step two (contract reading and clause annotation, Section 3.1.).

---

[25] Important caveat: The recall is 1 *to the best of our knowledge*. In other words, the software annotates all the potentially unfair clauses which were earlier or later detected by a human reader when annotating the corpus. This is not say that all the readers will agree on the totality of choices, nor that some of the clauses which nor that some of the clauses, if annotated by others, would not have been omitted. All this is due to the necessarily evaluative character of the exercise, the problem addressed in Section 2.2.

3.1.Machine Learning Approaches to the Construction of Datasets

As was explained in Section 2.2., the application of machine learning to the detection of potentially unfair clauses will be possible once two important preconditions are met. Firstly, a large enough dataset needs to be created; secondly, there needs to be a consensus among experts on what makes a good dataset. The important questions are: "Where would such a dataset come from?" and "among whom should consensus be achieved?"

Regarding the first question, the dataset could be constructed either by a large research group devoted to the text full time (one needs to remember that "large enough" in this context means encompassing at least several hundred, if not thousands, of documents), by the consumer protection organisation itself, or through some sort of community collaboration. For decades, machine learning research has been using community-contributed datasets, such as the UCI Dataset repository[26], KDD-Cup competitions[27], and others. The purpose of the existence of those datasets is to promote research in the area. In domains other than machine learning, there are still some available data that can be used for empirical analysis or prediction. Examples include Aletras et al. (2016), which used data available from the database of the European Court of Human Rights[28], and (Katz et al., 2014) which used the infamous Spaeth dataset[29] to predict the voting behaviour of judges in the U.S. Supreme Court. In the majority of cases though, the absence of research-specific data has led the researchers to manually create their own dataset that has allowed them to test their hypotheses, in a process widely known as *coding.* This is the problem one faces here.

Regarding a set construction by a consumer body, one could imagine that the lawyers, in conducting their tasks as usual, would also produce, as a sort of positive externality, annotated documents to be fed into the database. The value of this approach is that it wouldn't create much additional work, however, progress would arguably be slow. This is not to say that such an approach cannot be used as complementary to other ones. Another option is a community-based construction. Imagine that whenever a consumer body or organization, governmental or non-governmental, or a court, or researcher, or a legally-conscious user found a clause that is potentially

---

[26] https://archive.ics.uci.edu/ml/datasets.html

[27] http://aka.ms/academicgraph

[28] http://hudoc.echr.coe.int

[29] http://scdb.wustl.edu/

unfair in some contract, they allowed the machine to learn about it. Imagine that the set were stored in a cloud, and that hundreds of people were allowed to enrich it. Suddenly, the task of teaching the computer would be undertaken not by one, or five, but by hundreds of people. This could result in a database of immense size. Ideally, a pilot project involving a consumer protection organisation and a research group as a hub, collaborating with other networks, could be established.

Again, technicalities regarding how this could be conducted – whether anyone or only authorized bodies should have access, whether this would be voluntary or obligatory on the side of public bodies, for example, how to test the software before making the database usable etc. – may differ, but can be debated at the right time. What matters is that it is possible, even now, and could be done if only consumer protection organizations and bodies had the will. Whatever the answers to these questions, the construction of the dataset will not happen by itself – it must be facilitated in one way or another by policy makers. How exactly this should be done is the subject of another paper, and is a potential follow-up to this research.

3.2.A Web-Crawler: Automation of Documents' Collection and Communication

> web crawler (web spider)
> "An application that systematically and continuously traverses the World Wide Web
> for some purpose. Web search engines, for example, use web crawlers to refresh
> their data by visiting every website periodically. Performing such traversals is called
> spidering." (Butterfield & Ngondi, 2016)

Each and every reader is familiar with the technology of web crawlers, even if they have never heard the term before. That is because web crawlers are already extensively used by business. The most popular types of services that use this technology are search engines like Google, Yahoo or Bing. When one types a search phrase in the Google search bar, Google does not scan the internet looking for web pages, it scans its own database created, in advance, by web crawlers. What Google does is use this type of software to automatically scan and index websites, and to create a database containing hyperlinks leading to them, matched with saved content IDs, key words, categorizations etc. Google and every search engine repeat the same process quite often to keep their content up-to-date. Google's database is extremely sophisticated, and so is their web-crawling technology. However, for the purposes of enforcement of unfair terms rules, a much simpler technology would suffice, one that can already be designed. The only reason why consumer

protection organization/authorities are not using one yet is because they still have not seen proof of the concept to understand its potential.

As of now, we have automated the reading of terms of service and the detection of potentially unfair clauses, but a consumer protection authority/organization would still choose what terms of what service should be scanned. They would still need to decide to check an agreement of a particular service/platform, either by importing it straight from the website or by saving it to a document. This still requires time and knowledge about a service's existence. But it does not have to be this way.

It is perfectly possible to design a web crawler that would automatically traverse a part of the web in search of online services, look for their terms of service, import, save, and scan them and produce a report. Consequently, it is possible for consumer protection authorities/organizations to independently construct a database containing all the terms of use of online services out there, or in a particular market (this can be designed by the institutions themselves), with an indication of the potentially unfair clauses in the documents. This would allow the automation of step one.

Potential uses of such technology are immense. For example, the firm and authority could be automatically notified. One could imagine a crawler that would send an email to a company when it detects potentially unfair clauses in the terms of service. Such an email could contain a report of which clauses are potentially unfair in the contract, as well as pre-prepared guidelines on what unfair terms are and how to fix them. The email would inform the business that potentially unfair clauses were detected in their terms, inform them that the consumer protection authority/organization had also been notified and sent the report, and propose a way to proceed. For example, it could give the business some time (one month, two months) to amend their terms of service and request that the consumer authority be notified. If the company didn't comply, the authority would initiate legal proceedings against it. This would be the automation of step four. If this works, everyone can do it. The existence of such technology could empower many more legal actors to exercise abstract control of the fairness of terms. Only some parties have legal standing in court. Technology, on the other hand, is open to everybody, together with other tools of exercising pressure on business, like sending letters, launching consumer campaigns, etc. This, obviously, must take place within the boundaries of the rule of law (Brownsword, 2016).

Again, numerous technical questions arise: Would this be possible in the current state of law, or should the law be amended? Should there be a reverse burden of proof in the case of automatic notification, as is the case with the so-called "grey list" clauses? What should the time spans be? What would the legal status of the guidelines be? These are all questions to be addressed

in another paper, and will certainly require a lot of discussion. Initiating the discussion is one of the ambitions of this piece.

Conclusion

In this contribution the authors have argued that several tasks which currently need to be performed by humans in the process of abstract control of fairness of clauses can be automated. The theoretical argument for this claim has been offered, as has empirical evidence that at least one of these tasks, i.e., reading and annotating the terms of online services can already be automated here and now. However, the results of the research presented here are only the first step.

It is already possible to significantly speed up the detection process of unfair contractual clauses. If more is invested in the research like this, an epochal change can occur for consumer organizations. The software proposed in this paper is a useful tool, but only the beginning. It can be further developed in a number of ways.

What is relevant from a policy perspective is to convey the message that consumers can be better protected, and the resources of consumer authorities better used, if research continues on this path. There is no reason not to do it. Obviously, much more work remains to be done. Closer collaboration between legal scholars, engineers and consumer protection people is needed. Where else can this technology be used? What other problems do consumer bodies face? How can they be smoothened by the automation of enforcement? These questions can only be answered through a mutual dialogue between those fighting at the battlefield and experiencing the difficulties (consumer protection organizations and bodies), those who know the rules of the game (legal scholars) and those who know how to build the arms (engineers).

We can literally revolutionize the enforcement of the law regarding unfair contractual terms and we hope the research presented here is one of the sparks that will start the revolution.

Bibliography

Aletras, N., Tsarapatsanis, D., Preoţiuc-Pietro, D., & Lampos, V. (2016). Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective. *PeerJ Computer Science, 2*, e93. https://doi.org/10.7717/peerj-cs.93

Alpaydin, E. (2014). *Introduction to machine learning* (Third edition.). The MIT Press.

Alpaydin, E. (2016). *Machine Learning*. Cambridge MA: MIT Press.

Benkler, Y. (2006). *The wealth of networks : how social production transforms markets and freedom*. Yale University Press.

Boer, A. (2009). *Legal theory, sources of law, and the semantic web*. Amsterdam, Netherlands ; Washington, DC: IOS Press.

Branting, L. K. (2017). Data-centric and logic-based models for automated legal problem solving. *Artificial Intelligence and Law*, *25*(1), 5–27. https://doi.org/10.1007/s10506-017-9193-x

Brownsword, R. (2008). *Rights, regulation, and the technological revolution*. Oxford ; New York: Oxford University Press.

Brownsword, R. (2016). Technological management and the Rule of Law. *Law, Innovation and Technology*, *8*(1), 100–140. https://doi.org/10.1080/17579961.2016.1161891

Butterfield, A., & Ngondi, G. E. (Eds.). (2016). *A Dictionary of Computer Science* (7th ed.). Oxford University Press.

Chopra, S., & White, L. F. (2011). *A Legal Theory for Autonomous Artificial Agents*. Ann Arbor: The University of Michigan Press.

Chu-Carroll, J., Fan, J., Boguraev, B. K., Carmel, D., Sheinwald, D., & Welty, C. (2012). Finding needles in the haystack: Search and candidate generation. *IBM Journal of Research and Development*, *56*(3.4), 6:1–6:12. https://doi.org/10.1147/JRD.2012.2186682

Chu-Carroll, J., Fan, J., Schlaefer, N., & Zadrozny, W. (2012). Textual resource acquisition and engineering. *IBM Journal of Research and Development*, *56*(3.4), 4:1–4:11. https://doi.org/10.1147/JRD.2012.2185901

Costa-jussà, M. R., & Fonollosa, J. A. R. (2015). Latest trends in hybrid machine translation and its applications. *Computer Speech & Language*, *32*(1), 3–10. https://doi.org/10.1016/j.csl.2014.11.001

De Franceschi, A. (2016). *European contract law and the digital single market : the implications of the digital revolution*. Intersentia.

Goldsmith, J. L., & Wu, T. (2006). *Who controls the Internet? : illusions of a borderless world*. Oxford University Press.

Gowers, T. (2002). *Mathematics : a very short introduction*. Oxford University Press.

Grundmann, S., & Kull, I. (Eds.). (2017). *European Contract Law in the Digital Age*. Antwerp/Cambridge: Interesentia.

Han, J., Kamber, M., & Pei, J. (2011). *Data Mining: Concepts and Techniques, Third Edition* (3 edition). Haryana, India; Burlington, MA: Morgan Kaufmann.

Johnson, D. R., & Post, D. (1996). Law and Borders - The Rise of Law in Cyberspace. *Stanford Law Review*, *48*(5), 36.

Katz, D. M., Bommarito II, M. J., & Blackman, J. (2014). Predicting the Behavior of the Supreme Court of the United States: A General Approach. *arXiv:1407.6333 [physics]*. Retrieved from http://arxiv.org/abs/1407.6333

Lessig, L. (1999). Law of the Horse: What Cyber Law Might Teach, The. *Harvard Law Review*, *113*, 501.

Lessig, L. (2006). *Code version 2.0*. New York: Basic Books.

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. https://doi.org/10.1007/s10603-015-9303-7

Marr, D. (2010). *Vision. A Computational Investigation into the Human Representation and Processing of Visual Information*. Cambridge MA: MIT Press.

Micklitz, H.-W. (2010). Reforming European Unfair Terms Legislation in Consumer Contracts. *European Review of Contract Law*, *6*(4), 347–383. https://doi.org/10.1515/ercl.2010.347

Micklitz, H.-W., & Kas, B. (2014). Overview of cases before the CJEU on European Consumer Contract Law (2008–2013) – Part I. *European Review of Contract Law*, *10*(1), 1–63. https://doi.org/10.1515/ercl-2014-0001

Micklitz, H.-W., & Reich, N. (2014). The Court and Sleeping Beauty: The revival of the Unfair Contract Terms Directive (UCTD). *Common Market Law Review*, *51*(3), 771–808.

Micklitz, H.-W., Reich, N., Rott, P., & Tonner, K. (2014). *European Consumer Law: Second edition*. Intersentia.

Nebbia, P. (2007). *Unfair contract terms in European law : a study in comparative and EC law*. Hart.

Palfrey, J. G., & Gasser, U. (2008). *Born digital: understanding the first generation of digital natives*. New York: Basic Books.

Palka, P. (2017). Beyond Contract Law in the Regulation of Online Platforms: Terms of Service are not Contracts. In S. Grundmann & I. Kull (Eds.), *European Contract Law in the Digital Age*. Antwerp/Cambridge: Interesentia.

Russell, S. J. (Stuart J., & Norvig, P. (2014). *Artificial intelligence : a modern approach* (Pearson new international edition.). Pearson.

Sartor, G. (2011). Legislative Information and the Web. In G. Sartor, M. Palmirani, E. Francesconi, & M. A. Biasiotti (Eds.), *Legislative XML for the Semantic Web*. Springer.

Schulte-Nölke, H., Twigg-Flesner, C., & Ebers, M. (Eds.). (2008). *EC consumer law compendium: the consumer acquis and its transposition in the member state [i.e. states]*. Munich: Sellier. European Law Publishers.

Schulze, R., & Staudenmayer, D. (2016). *Digital revolution : challenges for contract law in practice*. Nomos/Hart.

Surden, H. (2014). Machine Learning and Law Essay. *Washington Law Review*, *89*, 87–116.

Trottier, D., & Fuchs, C. (2015). *Social media, politics and the state: protests, revolutions, riots, crime and policing in the age of Facebook, Twitter and YouTube*. New York: Routledge.

Uzuner, Ö., Zhang, X., & Sibanda, T. (2009). Machine Learning and Rule-based Approaches to Assertion Classification. *Journal of the American Medical Informatics Association*, *16*(1), 109–115. https://doi.org/10.1197/jamia.M2950

Wendehorst, C. (2016). Verbraucherrelevante Problemstellungen zu Besitz- und Eigentumsverhältnissen beim Internet der Dinge. *Studien Und Gutachten Im Auftrag Des Sachverständigenrats Für Verbraucherfragen Dezember 2016*. Retrieved from http://www.svr-verbraucherfragen.de/wp-content/uploads/2016/11/Wendehorst-Gutachten.pdf

Zittrain, J. (2008). *The Future of the Internet. And How to Stop It*. New Haven & London: Yale University Press.