The Issue of Proxies and Choice Architectures

Why EU law matters for recommender systems

Mireille Hildebrandt

Abstract

Recommendations are meant to increase sales or ad revenue, as these are the first priority of those who pay for them. As recommender systems match their recommendations with inferred preferences, we should not be surprised if the algorithm optimises for lucrative preferences and thus co-produces the preferences they mine. This relates to the well-known problem of feedback loops, filter bubbles and echo chambers. In this article, I discuss the implications of the fact that computing systems necessarily work with proxies when inferring recommendations and raise a number of questions about whether recommender systems actually do what they are claimed to do, while also analysing the often-perverse economic incentive structures that have a major impact on relevant design decisions. Finally, I will explain how the choice architectures for data controllers and providers of AI systems as foreseen in the EU's General Data Protection Regulation (GDPR), the proposed EU Digital Services Act (DSA) and the proposed EU AI Act will help to break through various vicious circles, by constraining how people may be targeted (GDPR, DSA) and by requiring documented evidence of the robustness, resilience, reliability and the responsible design and deployment of high-risk recommender systems (AI Act).

Keywords: micro-targeting, machine learning, behavioural profiling, political economy, behaviourism, Goodhart effect, affordance, capability

The Issue of Proxies and Choice Architectures	1
<i>Introduction</i>	3
<i>Defining RSs</i>	5
<i>A political economy of recommender systems</i>	6
<i>The business and the math of persuasion</i>	7
Hidden persuaders old and new	7
Design choices: the issue of proxies	8
<i>Behaviourist assumptions of micro-targeting</i>	9
Behaviourism old and new: the inversion of the proxy relation	9
Dark patterns	10
<i>Intermezzo on counter profiling</i>	11
<i>Human-machine feedback loops and the Goodhart effect</i>	13
Illusions of control over human behaviour: the return of the proxy	13
Goodhart effect	14
Adaptive anticipation	15
Methodological integrity in science and society	16
<i>The legal framework as a choice architecture</i>	17
Hijacking the concept of 'choice architecture'	17
The choice architecture for controllers under the GDPR	18
The choice architecture for providers under the AI Act	22
<i>Conclusions: deepening the Brussels effect?</i>	24

Introduction

Recommender systems (RSs) based on collaborative filtering inevitably create feedback loops, echo chambers and filter bubbles, because algorithms cannot be trained on future data. On top of that, the political economy that drives the incentive structure for providers and users of these systems creates perverse incentives, skewing the inferences in a direction that is favourable for those hoping to make a profit or win or confuse public opinion. Both issues are exacerbated by the fact that computing systems necessarily work with proxies when determining the target outcome (relevant recommendations) and when deciding on relevant feature variables (based on the distribution in training data).

This article puts forward two reasons why the issue of proxies is difficult to resolve. First, to resolve means first to accept that any inferred recommendation will always be limited by the choice of the proxies; the issue cannot be resolved by selecting perfect proxies that are identical with what they aim to map. Second, the issue can be mitigated by selecting better proxies, but only if 'better' relates to better attention paid to perverse economic incentives and to more acuity as to the disturbingly naïve behaviourist assumptions that plague the research design of collaborative filtering based on behavioural data. The first requires more research into the political economy that informs the research design of RSs, the second requires more research into the way the output of RSs is presented, notably the use of performance metrics that glean over the nature of computable proxies and their limitations.

The current and upcoming EU legislative framework creates new choice architectures for users (deployers) and providers of RSs, such that end-users' agency is taken seriously, and their capabilities are enhanced rather than captured in loops, bubbles and echo chambers. At the end of the day, these new choice architectures will not only constrain how RSs can be designed and developed; they will instigate a kind of 'by design' protection that is based on carefully selected checks and balances, meant to create a computational infrastructure that is not poised towards manipulation but towards informed interaction. This will require upfront involvement of the developers of these systems, who need to understand that law is not a boring bag of rules meant to constrain them but a set of dedicated incentives to build more resilient, robust, reliable and responsible RSs.

In this paper, I will focus on RSs that combine the objective of providing recommendations in whatever context with a business model that aims to persuade the end-user to prioritise and/or develop specific preferences, for instance resulting in long term user engagement [1] or in selection of products from advertisers that pay a higher conversion fee [2]. This is not to suggest that systems not geared towards such persuasion are neutral or objective – relevance always depends on purpose, context and agent; in that sense even a well-designed and properly deployed system is biased by definition, and this is not a bug but a feature [3].

Those eager to check out the potential impact of the EU legislative framework can start with the sections 'the intermezzo on counter profiling' (for the relevance of the

Digital services act or DSA)¹ and 'the legal framework as a choice architecture' (for the relevance of the General data protection regulation or GDPR² and the proposed AI Act³). They should, however, take into account that the article is not merely a legal analysis but aims to situate how and why law matters here. This means the legal analysis is focused on how the GDPR and the proposed DSA and AI Act re-configure the choice architecture of big tech in relation to the risks that RSs pose for human rights. The key point of the article is to better understand the issues of proxies that infest RSs and the nature of law as a means to protect individual human beings against detrimental effects of the pseudo-science that underpins many of the currently marketed RSs, notably where such effects are skewed against those already disadvantaged.

In the section 'defining RSs', I will put my cards on the table as to how I understand the role and function of RSs, emphasising that they constitute an inevitable instrument since we moved from information scarcity to information overflow. As Dewey explained, instruments are not neutral [4]. They can be developed and used in different ways that will reconfigure the goals they aim to achieve, even as the goals co-determine the instruments to achieve them. This also goes for RSs. In the section 'a political economy of recommender systems', I will explain how law determines the shape and the affordances of economic markets and how these markets define the power relations between different players in the RS ecosystem to assess what this means for the implied incentive structure. After thus sensitising the reader to the impact of economic incentives I will discuss 'the business and the math of persuasion', starting with Packard's seminal description of 'motivational research' in the 1970s, following up with current persuasion techniques and the design choices they involve. This will feed into a brief recall of the 'behaviourist assumptions of micro-targeting', as discussed in other work [5], focusing on old and new types of behaviourism and how they give rise to dark patterns [6]. This relates to the issue of proxies, i.e. the machine-readable variables that stand for relevant features and targets.

Instead of reinforcing the belief that micro-targeting works as claimed, I will pause for an 'intermezzo on counter profiling', a technique that aims to infer how RSs apparently profile their end-users, thus enabling profiling of the profilers. I then discuss 'human-machine feedback loops and the Goodhart effect', inquiring into the pitfalls of cybernetic control over human behaviour, explaining the notion of the Goodhart effect and how this relates to the kind of adaptive anticipation that is key to human interaction. Thus, I will trace the consequences of the Goodhart effect for both science and society. I conclude that developers of RSs must face the music and

¹ Proposal for a Regulation of the European Parliament and of the Council on a Single Market For Digital Services (Digital Services Act) and amending Directive 2000/31/EC, COM/2020/825 final, <https://eur-lex.europa.eu/legal-content/en/TXT/?uri=COM%3A2020%3A825%3AFIN>.

² 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), <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

³ Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial intelligence act) and amending certain union legislative acts, COM/2021/206 final, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>.

wake up from their behaviourist dreams, if they want to remain relevant in a world that is swiftly moving towards lifting the veil from unsubstantiated claims.

In the section on 'the legal framework as a choice architecture' I will turn the tables on the usual narratives about 'choice architecture' that focus on the choice architectures that RSs offer to enable nudging citizens and consumers. Instead, I will frame dedicated legislation as a choice architecture for those who provide or deploy RSs. I will explain how the GDPR and the AI Act build specific choice architectures for *controllers*, i.e. those who determine the purpose and means for processing personal data, and for *providers* of high-risk AI systems, i.e. those who develop (or have others develop) AI systems that they put on the market or put into service under their own name or brand.

In my conclusions, I will argue that both the GDPR and the AI Act will shape global economic markets in ways that will affect the design, default settings and interactive features of RSs. The type of constraints this entails will not obstruct research and innovation in RSs but instead foster and enable resilient, robust, reliable and responsible RSs where end-users' agency is respected, based on user participation rather than relying on reductive user modelling. This is the most salient way to face the issue of proxies where it comes to human behaviour.

Defining RSs

As any handbook will tell [7], [8] RSs are meant to offer relevant information in myriad settings, basically dealing with situations where an abundance of information makes it hard to detect and select what is required in a specific situation, or with situations where information itself cannot be accessed. Recommender systems are now part of search engines (the PageRank algorithm), streaming providers (music, movies, podcasts), social networks (rankings in timelines, news feeds), webshops (suggesting similar products or services), platforms that mediate the gig economy (recommending short stays or car rides) and will no doubt feed into cyberphysical infrastructures (internet of things [9], smart cities [10], connected cars, smart energy grids [11]). On top of that behavioural advertising, which runs the business model of 'free' services such as search engines and social networks, is a dedicated type of recommender system [12]. On the one hand, behavioural advertising recommends products and services to potential consumers, on the other hand, it recommends advertising space to potential advertisers. In the current global marketplace, many recommender systems are in the business of persuasion, often combining advertising or marketing with offering relevant information.

It is crucially important that we acknowledge that without RSs neither individuals nor private or public organisations would be able to navigate, let alone address the current information overload. This, however, does not imply that anything goes, nor that we can simply refer to an agreed-upon objective standard of what information should be presented to whom, by whom, when, how and in what context.

Technology is neither bad nor good, but never neutral [13]. In the next section, we will explore the incentive structures that currently drive the provision and deployment of RSs. Before that, we must clarify the backend of these systems and explain how its design informs their deployment and their impact on those targeted with their output.

RSs are based on a relatively small set of filtering mechanisms that are tuned to specific types of users, content or both. These mechanisms involve inferences from behavioural data or input explicitly provided by users concerning their preferences. The main filtering techniques are collaborative based (inferences based on aggregate user behaviour), content-based (retrieval of relevant documents in the context of a specific domain) or knowledge-based (often constraint-based, with input provided by domain experts) and various types of hybrids. This implies overlap with disciplines such as information retrieval, machine learning, symbolic reasoning, knowledge management and search. As this special issue is focused on 'human-machine feedback loops' we are reminded that human-machine-interfaces or human-machine-interaction (HMI) is another key discipline in this domain [14].

Attention to human-machine interaction is, however, often limited to end-user-friendly design and user modelling, hoping to achieve some goal as to the user's behaviour. However, for RSs to operate in efficient and effective ways, keen attention must be paid to (1) the interaction between those designing and building RSs and the system under development, as well as to (2) the interaction between those who deploy the system to offer specific services and the system they use, and, finally, to (3) the interaction between end-users and the RS they engage with. Interactionist approaches that focus on the third type of interaction (and maybe the second), overlook that the first and the second actually determine the affordances of the system that end-users get to interact with; together, they define the *choice architecture* [15] that controls what types of options the end-user has, thus constraining the kind of recommendations they will obtain. This invites a further inquiry into the economic incentive structures that keep developers, providers and those who deploy RSs on their toes.

A political economy of recommender systems

In 2008 we published *Profiling the European Citizen* [16], co-authored by computer scientists, lawyers, social scientists and philosophers. We focused on the computational, legal, social and underlying epistemological issues of using what was then called 'knowledge discovery in databases' (KDD), inspired by early work on the complexities and potential impact of 'group profiling' [17][18]. *Profiling the European Citizen* is still highly relevant in terms of issues such as bias, discrimination, privacy and automation bias, but also pivotal as to a proper understanding of the manipulative impact of design and deployment of KDD and similar systems (such as RSs). The volume engaged a relational understanding of knowledge, as emphasised in the follow-up work of 2018, subtitled '*cogitas ergo sum*' [19], acknowledging that the way others frame us is part of what shapes us. Recently, we have seen a number of best-selling works that make similar points about the impact of machine learning available in a more narrative style for a larger audience [20]–[23], including the oft-cited work by Zuboff on 'surveillance capitalism' [24]. My goal in this article is not to repeat these points made, but to explain how law shapes economic markets [25]–[27] and how that relates to the issue of the proxies that RSs engage to make their models work. In other words, my goal is to show the importance and impact of the choice of proxies when developing RSs and to clarify why law matters for getting things rights.

Law does not grow like grass, though it requires cultivation. Whereas grass is an organism shaped by myriad biological interactions, law is an *institutional fact* [28]–[30] largely shaped by written legal *speech acts* whose performative effects ‘make’ the law [31]. Being text-driven and based on natural language ensures a combination of semantic continuity and adaptive flexibility that is core to law and the rule of law [32]. The relationship between speech acts and the institutional facts they create is not causal but constitutional and the same goes for the institutional facts created by the law, such as economic markets. Law shapes the choice architecture offered by economic markets, both where it enables (contract law, tort law, property law) and where it constrains (private law, public law, human rights law). If contract law did not cover the legal effect of a lack of performance, we could never be sure about the future consequences of a contract, which might then come to depend on prevailing power relationships rather than on what parties agreed upon. This would dissolve the binding character of law that sustains economic markets. If property law would not provide right holders with broad powers to use, alienate or even destroy their property, markets-as-we-know-them would not exist. If private law did not void contracts based on undue influence or fraud, markets would soon become dysfunctional or collapse into patronage systems. Those who care about democracy and the rule of law should pay keen attention to how markets are constituted, how they distribute powers and liabilities and thereby both wealth and inequality [26].

Thus, economic markets do not grow like grass, but depend on how we create and sustain them by way of binding legal norms that in turn enable people to engage in binding agreements. This invites an inquiry into the kind of incentive structure that these markets themselves create, based on what complex interplay of economic power relationships is enabled. This is the objective of sketching the political economy of recommender systems, i.e. figuring out what drives their provision and use; what players are enabled or kept in the dark.

Once we can trace the economic power relationships, we can inquire into the way they influence the methodological integrity of these systems and their reliability in terms of the functionalities claimed for them. This will also allow us to test how the GDPR and the proposed AI Act may contribute to an incentive structure that favours dependable systems and responsible deployment. Different sets of legal norms create different types of markets; economic markets are not given but constituted, they can be reconstituted if we want to.

The business and the math of persuasion

Hidden persuaders old and new

To understand the political economy of recommender systems we must focus on the power relationships between different players. These are (1) providers of products or services, e.g. advertisers, energy suppliers, webshops, (2) publishers that host recommendations, e.g. websites, search engines and social media that sell space for advertising, (3) intermediaries that negotiate between publishers and providers and/or between publishers and consumers and (4) consumers or end-users targeted with supposedly or hopefully relevant recommendations.

In the context of the political economy of recommender systems the question of relevance is crucial. Relevance is a key concept in information retrieval [33], [34] and in the design of recommendation systems [7][8]. The key question, however, is relevant-for-whom: for the consumer or end-user? for those paying the bill, e.g. an advertiser? or for those making a profit, such as the publisher and the intermediary? In a market that thrives on free services it makes sense to check who pays what bill in exchange for what service, and even in a market where consumers do get to pay the bill (subscriptions) it makes sense to figure out what drives the recommendation (user preferences or user nudging [1]).

In 1957 Vance Packard described the cynical use of pseudo-scientific insights in marketing and advertising, in his seminal *Hidden Persuaders* [35]. In this work he revealed to what extent advertising and marketing were based on defunct 'insights' of commercially operating psychoanalysts, behaviourists and other folk, 'sold' under the heading of 'motivational research' (MR). MR was mostly grounded in unsubstantiated claims about how to manipulate the irrational unconscious of consumers or citizens (voters). In fact, MR was mainly deployed to persuade advertisers to buy into pseudo-scientific claims about how to persuade people to purchase what they may neither need nor want.

The conflation of machine learning and nudge theory that is the focus of this paper, seems a direct heir of MR, based on similarly untested claims regarding the effectiveness of behavioural targeting. And, once again, the success must be located in persuading the advertisers rather than voters or consumers. This, however, does not imply that the use of recommender systems makes no difference. As Packard notes in his final chapter on 'The Question of Morality', the aim to manipulate rather than inform people has moral implications:

What are the implications of all this persuasion in terms of our existing morality? What does it mean for the national morality to have so many powerfully influential people taking a manipulative attitude toward our society? Some of these persuaders, in their energetic endeavors to sway our actions, seem to fall unwittingly into the attitude that man exists to be manipulated.

Mapping the political economy of recommender systems aims to unearth the perverse incentives that drive systems meant to infer preferences. Actually, these systems are *rewarded* for influencing these preferences while catering to them. There is a loop here that keeps those targeted in an echo chamber, for instance *reinforcing* whatever elicits further 'user engagement' [1].

Design choices: the issue of proxies

The techniques of *collaborative filtering* that underly many currently deployed RSs assume several design choices that determine the output of RSs. First, the purpose must be specified as a task, which necessitates a formalisation that serves as a proxy for the 'real' purpose. The real purpose may be providing an end-user with what they like/need/want (already three very different concepts). Formalisation will probably be done by (1) stipulating that the system is given input data about what items/services/opinions they engaged with previously (where 'engaged' will require

formalisation in terms of e.g. click or purchasing behaviour) or (2) stipulating that the system is given input data about what items/services/opinions similar end-users engaged with (which entails formalisation of 'similar end-users'). The proxies involved are the measurable historical behavioural data of the end-user (e.g. their click- surf- or purchasing behaviour) and/or measurable historical behavioural data of similar end-users. Though it may be very interesting to make inferences and provide recommendations on this basis, one can debate whether an end-user's liking/needs/wants can be reduced to such input. Another way of framing this is asking whether the developers are doing a good job in 'groundtruthing' [36], i.e. in providing the learner algorithm with the right model for what would be a correct output, where that model is basically the distribution of the data on which the learner is trained, validated and tested. With an eye to methodological integrity, it is important to acknowledge that we have here a prime example of the use of a measure (the distribution of relevant data) as a target (steering end-users towards a similar distribution in the future). I will return to this under 'human-machine feedback loops and the Goodhart effect'.

Though proxies inevitably skew the output, this may have perverse consequences in the case of the RSs that are the target of this paper, as they will include proxies for persuasion-oriented goals that will inevitably influence the output (= recommendations). Again, almost by definition, the system will be using a measure (the distribution of relevant data) as a target (steering the end-user towards a similar distribution in the future). The point here is that, depending on the goals, 'the distribution of relevant data' will be compressed into different algorithms, resulting in different recommendations; 'the' distribution does not exist, unless in the form of replicating the entire data set [3].

We may conclude that big players who develop and/or provide RSs have a clear incentive to combine the purpose of providing relevant recommendations with persuasion purposes, meant to increase e.g. advertising revenue or to *prime for* specific types of political opinions, food choices or lifestyle behaviours.

Behaviourist assumptions of micro-targeting

Behaviourism old and new: the inversion of the proxy relation

Behaviourism is a way of 'doing' science that goes back to Pavlov, Skinner and Watson and informs both behavioural economics (nudge theory) and its 'reception' in computer science. In other work I have traced the problematic assumptions of behaviourism and their implications for the output of machine learning applications in learning analytics [5]. These assumptions regard the fact that behaviourism often inverts the relationship between concepts and their proxies, based on a belief that animal and human agency consists of discrete, observable, measurable units of behaviour (e.g. producing saliva, ticking boxes on a screen, 'firing of neurons') which are imprecisely expressed in concepts such as appetite, consent, thinking. This assumption is a bug from the perspective of scientific inquiry, because it confuses the modest proposition that we only have access to what we can observe with the

claim that only that which we can observe exists and/or matters.⁴ From the perspective of developers of RSs this – flawed - assumption may seem to be a feature, because it promises more precise and accurate knowledge due to the fact that supposedly vague concepts can be replaced by the behavioural primitives that inform them.

The problem of turning the relationship between proxies and concepts inside out was key to the so-called *Methodenstreit* that raged in economics and social science during the first half of the 20th century, with spill-over into the late 20th century 'science wars' [38]–[40]. The kind of behaviourism I am referring to underlies most work in econometrics and defines the point of departure of both Chicago Schools (rational choice theory as well as nudge theory). They both depend on what has been called 'methodological atomism' [41], a way of framing reality as an aggregate of discrete datapoints rather than e.g. acknowledging that a training dataset is nothing but a proxy for the 'truth' we want an algorithm to learn. The difference is that rational choice theory (classical economics) assumes that the atoms are rational and therefore predictable, whereas behavioural economics (nudge theory) assumes that the atoms are irrational but nevertheless predictable [42].

Micro targeting (MT) is core to the personalised recommender systems we discuss here, that is those intended to serve two masters: the end-user and the service provider. MT that is based on collaborative filtering requires hoarding of behavioural data of end-users (their own historical data and/or those of their similes). It is easier to 'sell' this with the hidden assumption that such data are the primitives that disclose our deepest intentions, desires and vulnerabilities. From a scientific perspective, however, it is important to acknowledge that these primitives – as well as any inferences based on them – are mere proxies for what the provider aims to know about the end-user. Those working with RSs based on MT are probably aware of this [43], but it is not in their interest to acknowledge this as that might deflate their business model.

Dark patterns

If the economic incentive structure invites business models that accept bugs as features, we should not be surprised that 'dark patterns' abound [6], aiming to sell RSs as innovative ways to influence end-users in a direction they might otherwise not have chosen. Though this may be qualified as manipulation, neoliberal ideology will present this as a way to grow the economy that is beneficial for those manipulated,⁵ while public policy pundits will portray this as benevolent paternalism [45]. Here the concept of a choice architecture is pivotal. 'Choice architecture' is a term of trade in nudge theory [15], referring to the types of choices that are offered to end-users, as consumers of goods, political opinion or whatever else. Choice architecture refers to a built environment (stones or software) that favours some choices over others, thus 'nudging' folk into the preferred behaviours of policy

⁴ Some versions of behaviourism avoid the reductive assumptions discussed here. Notably the work of G.H. Mead [37] takes a different perspective, informed by modesty and ingenuity rather than arrogance and intent to manipulate.

⁵ E.g. framing nudging techniques as 'tools that help people make better decisions "without forcing certain outcomes upon anyone" (Thaler et al., 2012, pp. 428–439)', thus declaring them beneficial by definition [44].

makers, webshops or advertising intermediaries. If it pays to lure consumers into dedicated behaviour patterns, preferably without their awareness, and if this is not prohibited, those with the means to do so may in some sense be forced to engage in such dark patterns as they believe they could otherwise be pushed out of the market. We should, therefore, not be surprised that 'dark patterns' of subliminal manipulation have indeed emerged.

What is more interesting is that some of the parties that play a key role in the political economy of RSs are pulling out, ignoring behavioural profiling in favour of e.g. contextual advertising. The New York Times stopped its behavioural advertising in Europe when the GDPR came into force and reported an increase in advertising revenue [46]. Before that, Proctor and Gamble substantially reduced its digital advertising budget, reporting a 10% increase in outreach one year later [47]. Doubts are being cast on the actual effectiveness of behavioural profiling [48], [49], papers are written on the limited gains for publishers [50] and on the pervasive misalignment of the economic and other incentives for different players [50]. My point here is not the moral and political argument against surveillance-based advertising [51], but a more foundational issue about the fault lines in the narrative about its effectiveness and reliability. Though it is crucial to make those moral and political arguments against manipulation, it is perhaps even more critical to better understand both the potential and the limits of micro targeting.

Intermezzo on counter profiling

In 2020, Facebook served researchers from the NYU Ad Observatory⁶ with a cease-and-desist letter, threatening legal action if they would not stop collecting targeting data regarding political ads [52][53]. Apart from closing the accounts of the researchers, Facebook also disabled application program interfaces (APIs) that allowed Facebook users to donate their data to the Ad Observatory to compare how people are actually targeted with political ads with the Facebook Ad Library that supposedly provides transparency about their targeting practices. The Ad Observatory found that many ads were missing from the Facebook Ad Library and tried to profile how Facebook's algorithms decide to target ads, something that Facebook does not disclose. In 2021, Facebook made changes to its website code that make it more difficult to identify sponsored posts, thus further frustrating efforts by independent researchers to scrutinize the way political advertising actually works on Facebook [54]. The Observatory basically engages in what I once coined counter-profiling, or profiling the profilers [55], [56].

In a report on behavioural biometric profiling from 2009, we distinguished between two types of transparency enhancing tools or TETs [55]:

Type A: legal and technological instruments that provide (a right of) access to (or information about) data processing, implying a transfer of knowledge from data controller to data subjects, and/or

⁶ See <https://adobservatory.org/>.

Type B: legal and technological instruments that (provide a right to) counter profile the smart environment in order to 'guess' how one's data match relevant group profiles that may affect one's risks and opportunities, implying that the observable and machine readable behaviour of one's environment provides enough information to anticipate the implications of one's behaviour.

Whereas type A depends on the willingness of a company or other entity to provide information about or access to relevant data (or on an effective right thereto), and on the technical ability to check whether this information is correct and/or the data complete, type B simply applies the same techniques of data-driven profiling that online and offline environments engage, to counter-profile that same environment. In other research, the DataBait tool⁷ was developed with a similar objective. Based on data-donations by users of social networks, we hoped to provide them with the kind of profiles that could be drawn about them, with a range of sensitive attributes [57]. Though the tool was not operational by the time the project finished, it might have generated similar attention from Facebook if it had been successful.

In 2021, Algorithm Watch announced that Facebook had threatened them with legal action if they continued their research on the Instagram Newsfeed algorithm, which Facebook found in violation of its Terms of Service (ToS) and the GDPR. Though both claims seem debatable if not far-fetched, Algorithm Watch decided to halt its research stating it does not have the means to start a legal battle with a tech giant. Note that the ToS prohibits scraping data from Facebook products with automated means, whereas Algorithm Watch only collected data from Facebook users who installed a plug-in to share what ads they receive. Note also that any content from other users was immediately and automatically removed, to be in compliance with data protection law (as these other users had not given their consent). Algorithm Watch then addressed an Open Letter to Members of the European Parliament, pressing for legislation that should enable this type of public interest research, more notably asserting the need to uphold and refine art. 31 of the proposed Digital Services Act (DSA)⁸ that gives 'vetted researchers' with academic affiliations access to platform data for public interest research.⁹

In point of fact art. 29 and 30 of the DSA require transparency regarding parameters and other relevant data whenever recommender systems and online advertising are used by very large online platforms. Moreover art. 29.1 DSA stipulates that the providers of these platforms must clarify in their terms of service:

any options for the recipients of the service to modify or influence those main parameters that they may have made available, including at least one option which is not based on profiling as defined in Article 4 (4) of Regulation (EU) 2016/679. Providers of very large online platforms shall

⁷ See <https://www.usemp.eu>.

⁸ Proposal for a Regulation of the European Parliament and of the Council on a Single Market For Digital Services (Digital Services Act) and amending Directive 2000/31/EC, draft as amended by the Council of the EU on 18 November 2021, <https://data.consilium.europa.eu/doc/document/ST-13203-2021-INIT/en/pdf>.

⁹ The Letter can be downloaded here: <https://algorithmwatch.org/en/defend-public-interest-research-on-platforms/>. Algorithm Watch requires that art. 31 should mention other sources of public interest research, notably civil society organisations and journalists. For a similar call to uphold and refine art. 31 see this Open Letter, signed by 50 organisations: <https://en.panoptikon.org/fix-algorithms>.

also make this information directly and easily accessible on the specific section of the online interface where the information is being prioritised according to the recommender system.

Counter profiling, however, is not equivalent with providers having to give users a chance to modify or influence a parameter. Rather, counter profiling enables (1) a better understanding of whether recommender systems deliver what they are claimed to deliver and (2) it can trace some of the actual consequences of microtargeting, especially at the aggregate level. It might enable us to demonstrate how the entire information ecosystem is distorted in the process of feeding targeted individuals humbug, even if those individuals have learnt to resist and avoid manipulation.¹⁰

Human-machine feedback loops and the Goodhart effect

Illusions of control over human behaviour: the return of the proxy

The interesting question is why Facebook gets so upset when their targeting data are scrutinized from the outside, that is from the perspective of their users – based on the same types of techniques they deploy to target them. The Ad Observatory found that the Facebook Ad Library (a repository of relevant data made available by Facebook) was incomplete in ways that contradict the claims made by Facebook. It seems the repository is carefully curated to prevent anybody from gaming the Facebook algorithms and/or from figuring out the real-world (in)effectiveness of their behavioural profiling. Returning to the current incentive structure (as shaped by relevant legal architectures), we could simply conclude that ‘the problem with Facebook is Facebook’, evidenced by e.g. the work of Vaidhyanathan [59] and inside testimonies that lift the veil of Facebook’s hidden agendas [60]. My aim here is, however, another one. I believe that more is at stake than dirty politics within big tech, enabled by inadequate legislation or spineless enforcement. The point is also that the underlying belief in our ability to manipulate human ‘behaviours’ (calculable changes in the state of human beings) is deeply flawed, mistaking proxies for what they are meant to represent.

To understand the issues at stake at a deeper – and simultaneously highly practical - level, we must pay keen attention to a crucial caveat that haunts all attempts to control human behaviour based on the measurement of discrete behaviours. Measurements of human interaction requires quantification, but quantification paradoxically requires prior qualification [61]. To measure ‘how many people like long hair’ one has to qualify certain behaviours as ‘liking’ and certain hair as long hair, which will involve all kinds of shortcuts or proxies. In times of Facebook, ‘liking’ may refer to clicking a thumbs-up icon, but that’s a very remote proxy for people actually liking something. Long hair may refer to certain dog breeds, to women or to men, to various types of hair (we do not only have hair on our heads, mind you). You may notice that specifying the kind of hair intended is less cumbersome than specifying what we mean with ‘liking’. Some concepts are more rich, depending on a

¹⁰ I would recommend (sic!) reading Harry Frankfurt’s seminal *On Bullshit* to better understand how the perverse incentives of the political economy of recommender systems destroys the conditions of possibility of democratic decision making [58].

more complex type of tacit knowledge [62], while such concepts are often what Gallie called 'essentially contested' [63], meaning that there is fundamental disagreement on what they refer to. Resolving that disagreement by singling out one meaning and imposing it on the research design to solve the engineering problem may result in skewed output. In short, whatever formalizable and quantifiable proxy one selects for 'liking', the quantification will always (1) require an act of qualification of concrete instances as 'liking' (whether that qualification is done by hand or automated) and (2) involve opting for one interpretation over others, thereby *loosing* the kind of tacit knowledge that grounds our understanding of 'liking' while *gaining* computability. In the case of supervised machine learning this is visible in the need to pay for a cohort of labourers to label (qualify) the training data, hiding the act of interpretation that is involved in this qualification.¹¹ On top of that, those who instruct the labourers impose their qualification on the training data, which will constrain the potential output of the learning algorithm as it can only see what has been labelled (framed) as 'liking'.

As indicated above, RSs that seek to offer what consumers 'like' require human-system interaction on three levels, that of (1) development, (2) deployment, and that of (3) micro-targeting. Interestingly, the interaction between developers and the system under construction will impose constraints on what the system is capable as recognising as 'liked' items (whether an opinion or a product or service). The same goes for the interaction between whoever deploys the system, as they will choose certain default settings, thus further shaping the choice architecture of the end-user (who will receive some recommendations at the cost of others). At the level of the design of the system certain behaviours will be used as a measure of what users like, and because that measure is then used as a target the targeted human will end up in a feedback loop that keeps them in an echo chamber or filter bubble. The filters that are created at the level of design and deployment turn into a rather boring menu of recycled preferences (the preferred term for 'likes'). Because human beings are whimsical, ingenious and forever anticipating how others anticipate them they will stop 'liking' what they have been foreseen to 'like' and behave in ways the system may not capture because it cannot see outside the qualifications that trained its algorithms (even if these qualifications were inferred based on unsupervised learning).

Goodhart effect

In economics this has been identified in 1975 by Goodhart as the adage that [66]:¹²

Any observed statistical regularity will tend to collapse once pressure is placed upon it for control purposes.

In later work by Strathern this has been summed up as 'the Goodhart effect':[68]

¹¹ On the human labour that goes into labelling training data see [64], [65].

¹² Social scientist Campbell had observed the same [67]: 'The more any quantitative social indicator is used for social decision-making, the more subject it will be to corruption pressures and the more apt it will be to distort and corrupt the social processes it is intended to monitor.'

When a measure becomes a target, it ceases to be a good measure.

The reason for this is simply that human agents will respond to the use of the measure with a change in behaviour that invalidates the measurement as such. This goes for the behaviour of those using the measure as a target, and the behaviour of those targeted. The measurement therefor keeps lagging behind what it aims to measure, while probably achieving numerous unintended side-effects. Strathern demonstrated this in her seminal article on the use of quantitative indicators for the quality of higher education, for instance, the number of students that enrol in a university, obtain their diploma, find work after graduating or, on the side of research, the number of publications, citations and journal impact. Obviously, once researchers are aware that their livelihood will depend on their measurable output they will change their expectations as to what counts as quality and change their behaviour to ensure a strong 'track record', while the management layer will change its assessment of quality research from one based on skilled intuition and content to one based on quantifiers and metadata. This is not a matter of researchers trying to game the system but rather a matter of them becoming part of a feedback loop that invites all kinds of 'elephant paths' to achieve the goals set for them (e.g. publishing the same content with slight tweaks in different journals, claiming novelty when merely repeating state of the art). This need not be a matter of deliberate deception but will rather be an intuitive change of direction that becomes the new common sense, resulting in new journals to accept the growing number of articles, infeasibility to read all that is published, in turn resulting in increased dependence on citation metrics. This is where the circle closes; the measure of high quality (high impact) becomes the target (publishing in high impact journals), driving a competitive 'market' for publications that define career paths. In other words: while the proxy of the target is mistaken for the target itself, the aim of high quality research turns into the aim to provide observable, quantifiable evidence of high quality. Again, the proxy is taken for what it stands for, inverting the relationship along the lines of good old behaviourism.

As we are restricting ourselves to systems that seek to *influence* what consumers 'like' while *mining* what they 'like', the Goodhart effect will play out in even more perverse ways than it would if the only target were to provide people with their inferred preferences. A retailer offering the right product to a customer may be successful, an advertiser creating a need/desire/want for a particular product may be equally effective, but trying to simultaneously create, predict and target individual preferences on a massive scale will not work in the long run. Not because recommender systems are not sufficiently sophisticated, e.g. based on deep learning techniques testing ever more parameters resulting in mathematically precise personalised recommendations. On the contrary, whatever mathematical precision is achieved, it will forever be off the mark as people intuitively anticipate how they are profiled, shift their expectations, and change their interactions. This also implies that 'fixing' the Goodhart effect by way of better metrics, as some suggest [69], may instead exacerbate rather than solve the problem.

Adaptive anticipation

The anticipation I refer to is not based on the kind of calculation that is the hallmark of rational choice theory, but rather on the intuitive learning capacity that some hold

to be 'irrational but predictable'. Instead of framing this intuition as irrational (wrongly biased), as Kahneman does [70], we can better understand this anticipation as based on our ability to deal with complexity by way of smart heuristics (rules of thumb) that are developed in tune with our situated or ecological rationality, as Gigerenzer argues [71].

Instead of framing our rationality or irrationality as properties of atomistic individuals, human rationality should be understood as a relational 'thing' that hinges on human interaction within a particular environment. Rational choice theory's assumption that human beings are optimisation machines hoping to maximise their own interests/preferences/benefits, rational choice theory is better understood as part of a utilitarian ideology that similarly underpins behavioural economics' nudge theory. The only difference between the two is the latter's caveat that humans suffer from cognitive biases that distort their ability to make the right choice, which is nevertheless the same type of rational choice that informs the other Chicago School. Based on a specific interpretation of Darwin, this ideological portrayal of individual human beings as intent on optimisation of their own interest at the cost of others has been 'naturalised' as a given biological fact [72], often invoking hunting and gathering societies as forever determining our competitive-aggressive nature.¹³ Human beings, however, are far too complex, situated and adaptive to be reduced to *homo economicus*, noting that culture, shared values and normative frameworks are dynamic and forever on the move.

The Goodhart effect, then, is not caused by folk trying to game the system but by interacting people that – while also interacting with RSs – dynamically intuit how they are being framed with what potential consequences; they will continuously reconfigure their responses accordingly. Such reconfiguration does not depend on them correctly intuiting how they have been framed but it does imply that the distribution of historical behavioural data will be different from that of future data. The most fundamental assumption of the kind of machine learning that drives many RSs, namely that these distributions are similar [3], [75], is wrong where it concerns human behaviour. In machine learning terms: historical data is not necessarily a good proxy for future data. Whoever build their castle on the proxy without taking this into account, are building on sand.

Methodological integrity in science and society

The bottom-line of the previous sections is that the combination of behaviourism, machine learning and nudge theory will not deliver what it promises while nevertheless creating havoc. The foundation of this toxic cocktail of (1) hidden behaviourist assumptions, (2) mistaking a proxy for the thing it supposedly stands for and (3) automating attempts to manipulate user engagement, reconfigures the private and the public sphere based on what is now euphemistically called 'misinformation'. From the perspective of computer science we thus face a detrimental lack of methodological integrity, due to the deployment of mathy pseudo-

¹³ The more interesting insights in the economic incentives in societies of hunters and gatherers can be found in the seminal work of anthropologist Sahlins [73], who also wrote saliently about attempts to ground human nature in naive interpretations of Darwin [74].

science [76], or – a more modest claim – mistaking exploratory research design for confirmed output [77], or – the internal critique – an incentive structure that invites ML papers that focus on algebra instead of new insights [78].

In the meantime, this is not just about the discipline of machine learning. This is about how we navigate the real world, because in the end that is what RSs aim to do: to improve our capabilities for navigating our physical and social environment in real life (IRL). To achieve methodological integrity within science in order to protect society against the deleterious effects of pseudo-scientific predictions we need to change human-machine-interaction at the level of the providers of RSs and of those who deploy them. This will require changing their choice architecture, to make sure the end-users get the recommendations they need/like/want in a way that respects their agency.

The legal framework as a choice architecture

Hijacking the concept of 'choice architecture'

Nudge theory has invented the concept of a choice architecture, to frame a deliberately designed environment that determines the types of choices available. The objective of creating a particular choice architecture is to exploit one or more cognitive biases to lure human agents into whatever behaviour is preferred by the provider. Nudge theory clarifies that RSs are designed to not only retrieve but also to influence end-user preferences. In the US, nudge theory and choice architecture have been embraced by public policy gurus such as Cass Sunstein, who argues that public policy should make use of them to help citizens to make the right choices [45]. This is qualified as libertarian paternalism, justified by the need to correct supposedly predictable irrational bias in human decision making. As discussed above, nudge theory is built on sand. I believe, however, that the concept of a choice architecture becomes very interesting if we liberate it from behaviourist assumptions and from attempts to influence people 'behind their back', instead of having them participate in the design of their environment. One could even say that this is what democracy is all about [79]: making sure those who suffer or enjoy the consequences of the built environment have a say in how this is done, while noting that 'built' refers to both our physical and our institutional environments.

Freed from naïve interpretations of human agency we can connect the concept of a 'choice architecture' with the concepts of 'affordance' and 'capability' that both highlight the relationship between a human agent and their environment. Together, affordance theory [80][81]¹⁴ and capability theory [83][84] understand agents in terms of what an environment 'affords' them and of what affordances they are 'capable' of acting upon. From that, ecological and relational, perspective, the notion of a choice architecture refers to what a particular built/designed/institutional environment affords specific types of agents in terms of the types of choices they can and cannot make, depending on their capabilities. This refurbishing of 'choice architecture'

¹⁴ Though I am aware of the prominence of Norman's [82] reconfiguration of Gibson's affordance theory in the context of human centric design, I am here referring to Gibson's original approach and e.g. that of [81], which provides for keen attention to the role of affordances in the constitution of human agency – and vice versa.

enables keen attention to how affordances depend on agents and vice versa – while also taking note that much of our environment is not built in stone but in terms of what has been called institutional facts.

A proper understanding of the notion of 'institutional facts' is key to understanding choice architectures in real life, as opposed to those being designed and tested in controlled environments. Institutional facts are facts created by performative speech acts; that is speech acts that 'do' what they 'say', where speaking is acting. The seminal example is the civil servant declaring a couple husband and wife, thus instituting their marriage, with all the legal effects attributed by the law of the land to the status of a marriage. The civil servant does not 'cause' them to be married, nor 'influenced' them into concluding the marriage, the declaration actually marries them.

Legal frameworks such as data protection law (GDPR)¹⁵ and the proposed regulation for AI systems (AI Act)¹⁶ create a choice architecture for controllers and processors who process personal data (GDPR) and for providers and users of AI systems (AI Act). This legal choice architecture affords them a dedicated set of choices on how to design, further develop and deploy their systems. The legislation aims for legal protection by design and default [85], thus requiring specific choice architectures to be engineered to empower end-users. The law itself, however, consists of institutional facts that are constituted by the written legal speech acts of the European legislature, articulated in natural language and thus explicit, multi-interpretable by default and therefor contestable. The purpose of the law is not to nudge folk behind their back into behaviours that legislatures prefer for them. The aim is to make the legal norms that a democracy has agreed upon explicit, enforceable and simultaneously. The law and the rule of law thereby respect and appeal to human agency in a way that persuasive RSs do not. This raises the fascinating question of what 'legal protection by design' could be.

The choice architecture for controllers under the GDPR

As discussed above, economic markets do not grow like grass, they are not brute facts but complex institutional facts with far reaching performative effects that co-determine the choice architecture faced by users and end-users of RSs such as corporations, consumers, employers and employees, house owners, tax payers, those who seek employment as well as government agencies, energy suppliers and big tech platforms.

The internal market of the EU is a prime example of a market that is created, enabled and constrained by private and public law of the member states and by myriad legal instruments enacted by the EU legislature and further developed by the Court of Justice of the EU (CJEU). More specifically, the choice architecture of those who process personal data is determined by the Charter of Fundamental Rights of

¹⁵ 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), <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

¹⁶ Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules of artificial intelligence (Artificial Intelligence Act) and amending certain union Legislative Acts COM/2021/206 final, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>.

the EU, the GDPR and other applicable legislative instruments, the relevant case law of the CJEU and other relevant sources of law (notably the Opinions of the European Data Protection Board and the European Data Protection Supervisor), together coined as the *EU data protection acquis*. For all corporations wishing to compete on the internal market of the EU, this *acquis* has relevance, even for those operating from outside the EU. This has been coined the Brussels effect [86] [87], highlighting how EU law shapes the global incentive structure for and of economic markets.

In this section I will select the most salient elements of the choice architecture of the EU data protection *acquis* insofar as relevant for RSs. The ultimate goal is to emphasise how this will also affect the design of RSs' backend, and transform the way developers, providers and users of RSs interact with them.

Current RSs are often based on the collection of online and possibly offline behavioural data (e.g. clickstream and surf behaviour, purchasing behaviour, location, traffic and navigation data, and data collected via quick response (QR) codes). The first thing to note here is that the data that are collected are considered personal data if the related person can 'reasonably likely' be identified. Due to the linkability of much behavioural data with other data this will often be the case [88],¹⁷ triggering the applicability of the GDPR. Once the GDPR applies, those who determine the purpose and means of processing become liable for complying with the GDPR, which will mostly be the user of the system (not being the end-user but e.g. a webshop or the provider of a social network or search engine). This entity is called 'the controller',¹⁸ and they have to ensure that they have a legal basis for the processing, which may be consent or a legitimate interest.¹⁹ Neither can be taken for granted, as the requirements for informed consent are high²⁰ and data subjects (those whose data is processed) must be able to withdraw their consent at any time in way that is as easy as when they gave their consent.²¹ The legitimate interest of the controller can only be used as a legal basis if it is not overruled by the interests, rights and freedoms of the data subject,²² noting the data subject has a right to object

¹⁷ Art. 4(1) GDPR: 'personal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person'.

¹⁸ Art. 4(7) GDPR: "controller" means the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data; where the purposes and means of such processing are determined by Union or Member State law, the controller or the specific criteria for its nomination may be provided for by Union or Member State law'.

¹⁹ There are actually six potential legal bases available in art. 6 of the GDPR: consent, contract, vital interests of the data subject or another, a legal obligation, performance of a task in the public interest or in the exercise of public authority, legitimate interest of the controller.

²⁰ Art. 4(11) GDPR: "consent" of the data subject means any freely given, specific, informed and unambiguous indication of the data subject's wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her'.

²¹ Art. 7(3) GDPR: 'The data subject shall have the right to withdraw his or her consent at any time. The withdrawal of consent shall not affect the lawfulness of processing based on consent before its withdrawal. Prior to giving consent, the data subject shall be informed thereof. It shall be as easy to withdraw as to give consent.'

²² Art. 6(1)f GDPR: 'processing is necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child.'

to the processing notably when based on the controller's legitimate interest.²³ This may sound dry and boring, reminding the reader of consent banners that pop up when surfing the web, disturbing the seamless 'user experience' of the end-user, in turn disrupting the seamless and subliminal influencing mechanisms that drive the business models of the users of RSs.²⁴ Though it is tempting to frame these legal requirements as obstructions, they may in fact enhance end-user agency. The sand that is thrown into the machine of smooth intuitive interactions *creates a different choice architecture* that forces all parties to pay keen attention to their different roles and reminds them of the fact that sharing data is a choice with consequences. As those consequences should be clear for end-users, cookie banners could be an example of what some have called 'apparency' [89] or 'actionable transparency', and to the extent that banners slow down both surfing and tracking it may be an example of what Paul Ohm called 'desirable inefficiency' [90]. In saying 'should be clear' and 'could be an example' I wish to highlight that the choice architecture of the GDPR has *force of law* not *force of technology*; it does not preclude disobedience and much will depend on (1) the extent to which legal norms are internalized as the right way to act and on (2) the way legal norms are enforced (noting (1) and (2) interact). The transparency requirements imposed by the GDPR nevertheless reconfigure the choice architecture of controllers who should be open as well as specific about their purposes to be lawful.²⁵ This entails that they should owe up to the fact that their recommendations have the specific purpose of increasing their own revenue, which will allow end-users to foresee potential bias that is inherent in the recommendation.

This brings us to the purpose limitation principle. The requirement of necessity that is core to the GDPR's purpose limitation principle is perhaps even more important than the legal basis. The purpose limitation principle restricts all processing to what is necessary for the specified purpose,²⁶ noting that the principle obliges controllers to define one or more legitimate purposes and to make them explicit. In line with that, consent is only valid if provided for a specific purpose and if the processing for which consent is given is necessary for that purpose.²⁷ This is related to the data minimization and storage limitation principles that *configure* the controller's *choice architecture* in a similar way: they are bound to choose processing operations that are not merely appropriate but necessary for the intended purposes. This implies

²³ Art. 21(1) GDPR: 'The data subject shall have the right to object, on grounds relating to his or her particular situation, at any time to processing of personal data concerning him or her which is based on point (e) or (f) of Article 6(1), including profiling based on those provisions. The controller shall no longer process the personal data unless the controller demonstrates compelling legitimate grounds for the processing which override the interests, rights and freedoms of the data subject or for the establishment, exercise or defence of legal claims.'

²⁴ Note that consent is also required by art. 5.3 of the ePrivacy Directive, even where it concerns the processing of data that do not qualify as personal data: 'Member States shall ensure that the storing of information, or the gaining of access to information already stored, in the terminal equipment of a subscriber or user is only allowed on condition that the subscriber or user concerned has given his or her consent, having been provided with clear and comprehensive information, in accordance with Directive 95/46/EC, inter alia, about the purposes of the processing. This shall not prevent any technical storage or access for the sole purpose of carrying out the transmission of a communication over an electronic communications network, or as strictly necessary in order for the provider of an information society service explicitly requested by the subscriber or user to provide the service.'

²⁵ Art. 12-14 impose transparency obligations on controllers, notably art. 13(1)c and 14(1)c: 'the controller shall provide the data subject with all of the following information: (...) the purposes of the processing for which the personal data are intended as well as the legal basis for the processing'.

²⁶ Art. 5(1)b GDPR: 'Personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes; (...) ('purpose limitation').

²⁷ Art. 6(1)a GDPR: 'the data subject has given consent to the processing of his or her personal data for one or more specific purposes'.

keen attention to what data is necessary to provide relevant recommendations and – as indicated above – it would require those deploying an RS to be explicit as well as specific about additional purposes that factually determine what and how data is processed or further processed. I could imagine that consent is freely given for the purpose of providing the end-user with relevant recommendations, whereas consent to provide recommendations that are foremost in the interest of the service provider, advertiser, political party or tech platform will most probably not be given, unless based on undue influence (e.g. dark patterns, or making the provision of a free service dependent on consent for additional processing).²⁸

There are more and other elements in the GDPR that co-determine the choice architecture that defines the market for RSs within EU jurisdiction, and thus also impacts design decisions for RSs that configure the backend system that end-users cannot access or control. For example, we can look into the prohibition of fully automated decisions that have legal or similarly significant effect on the end-user.²⁹ There is no case law on this particular element of the GDPR from the CJEU, but some relevant case law at the national level. The prohibition has three exceptions: consent, a legal obligation or contract. Those exceptions, however, are conditional upon certain safeguards, one of which is the obligation to provide meaningful information about the logic of processing³⁰ which has caused a flurry of research into what some have called 'explainable AI' [91][92]. Here again we see that *the choice architecture* for providers and users of RSs is reconfigured by such obligations, prompting providers to require that developers ensure that decisions taken by an RS are explainable to those targeted. One of the crucial questions here is what constitutes a 'decision' in the sense of the GDPR, taking note that both during the design and during the deployment of a RSs, many technical decisions are taken fully automatically by a system that may have a significant impact on end-users [93].

Law, economic markets and RSs interact in myriad ways. Considering the perverse effects of deploying RSs that are dependent on behavioural data (see above on the Goodhart effect) and the even more detrimental effects of optimizing recommendations in view of the interests of those who deploy the system rather than the end-users, we should be grateful for the choice architecture that is put forward by the EU data protection acquis, more specifically the GDPR. It is not so much an obstruction of seamless user experiences and innovative service provision, but instead an incentive to pay keen attention to the methodological integrity of developing RSs. Do they serve the claimed intended purpose? Is the processing of personal data necessary for that purpose? The GDPR contributes to 'good' innovation rather than innovation 'per se' and it may thus help to break the vicious

²⁸ Art. 7(4) GDPR: 'When assessing whether consent is freely given, utmost account shall be taken of whether, inter alia, the performance of a contract, including the provision of a service, is conditional on consent to the processing of personal data that is not necessary for the performance of that contract.'

²⁹ Art. 22(1) GDPR: 'The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.'

³⁰ Art. 15(1)h GDPR: 'The data subject shall have the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed, and, where that is the case, access to the personal data and the following information: the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.'

circle of perverse feedback loops. This will hopefully contribute to new ways of developing RSs, based on end-user *participation* instead of end-user *modelling*.

The choice architecture for providers under the AI Act

The EU AI Act (AIA) that was proposed in 2021 is still under consultation and negotiation while finalizing this paper. As the final text is not yet available, I will restrict myself to highlighting the groundbreaking nature of the Act, the choices made when defining AI and the kind of requirements facing providers of high risk AI systems. Though the current proposal is not perfect in any sense (see my feedback to the Commission [94], the analysis below should illustrate how the Act may reconfigure the choice architecture of the providers of what the Act calls 'AI systems'. As to its actual effectiveness much will depend on the budget allocation to national supervisors and on the upcoming revision of the EU Product Liability Directive which may impose strict liability on providers of high-risk RSs,³¹ especially if their Conformity Assessment turns out to be fake news.

The AI Act will have direct effect in all the Member States of the EU. It addresses the impact of AI systems on safety, health and fundamental rights and imposes a set of dedicated requirements on the providers of high risk AI systems. In the proposal, an AI system is defined in art. 3(1) of the AIA as

'software that is developed with one or more of the techniques and approaches listed in Annex I and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with'

The techniques and approaches listed in Annex I are:

- (a) Machine learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning;
- (b) Logic- and knowledge-based approaches, including knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning and expert systems;
- (c) Statistical approaches, Bayesian estimation, search and optimization methods.

Let's first note that the definition does not concern AI as a research domain or as an attempt to imitate or improve upon human intelligence. The definition targets systems and defines them in terms of four conditions: (1) it must be a software

³¹ Report from the Commission to the European Parliament, the Council and the European Economic and Social Committee. Report on the safety and liability implications of Artificial Intelligence, the Internet of Things and robotics, Brussels, 19.2.2020 COM(2020) 64 final, <https://eur-lex.europa.eu/legal-content/en/TXT/?qid=1593079180383&uri=CELEX%3A52020DC0064>.

system (so hardware without a software component or mere computer code are excluded), (2) the software system must have been developed by way of one or more of the techniques in the Annex, which includes both machine learning and logic- or knowledge based systems and various other techniques that seem already implied in the machine learning approaches, (3) the system is developed based on human defined objectives, which is inevitable but important as it links to the concept of intended purpose that is key to the Act and (4) to count as an AI system for the AIA it must generate outputs that influence the environments it interacts with. The latter means that an excel sheet in itself is not an AI system but might be so if it generates outputs that impact its environment, based on input from its environment. If an excel sheet is used to decide on social security benefits based on inputs about the targeted applicant, I could imagine this qualifies as an AI system. Actually this seems perfectly reasonable to me, the excel sheet is a logic based systems that performs certain operations on input data to generate an output. This output can be a recommendation and it should be clear that most RSs will squarely fall within the definition of AI systems in the current version of the Act.

The more interesting question is whether an RS is a high risk system, or rather under what conditions it could be qualified as such. This will depend on the impact they have.³² Annex III sums up a set of contexts and applications that involve high risk, e.g. when influencing access to education, employment, social benefits or creditworthiness, and various uses in public administration, including border control and criminal justice.³³ As to marketing and advertising the AIA prohibits the use of AI systems for subliminal manipulation and exploitation of vulnerable groups if such usage 'causes or is likely to cause (...) physical or psychological harm' for an individual person.³⁴ This may be hard to prove, but depending on how the burden of proof is distributed and on whether collective action is possible, I could imagine that certain types of RSs may at some point be prohibited.

If a RS were to be qualified as high risk, a whole range of dedicated legal obligations apply, mainly to the providers of these systems. They need to put in place a risk management system that assesses risks to health, safety and fundamental rights, both when the RS is used for its intended purpose and for reasonably foreseeable other purposes. For instance, the intended purpose of learning analytics may be to assess the progress of students, but once it is used to recommend exclusion the stakes become higher and non-discrimination could become a serious risk. The Act stipulates keen attention to the choice and curation of training data, and to various types of testing, while also emphasising proper performance metrics, robustness and cybersecurity, coupled with documentation and record keeping, automated logging and post market monitoring. The provider must put in place a quality management system and ensure both by design and by way of instruction that human oversight is meaningful, effective and practical.

In other words, the AIA restricts the kind of choices that providers of AI systems can make when leading the development of these systems, noting that the obligations

³² For the definition of high risk systems, see art. 6 AI Act which refers to Annex II (mainly focusing on safety and health risks) and Annex III (mainly focusing on risks to fundamental rights).

³³ Title III, chapter 2, art. 8-15 AI Act.

³⁴ Art. 4 AI Act.

often concern the backend system, the research design and the default settings. By thus skewing the choice architecture of providers, the AIA will favour resilient, robust, reliable and responsible (4R) AI systems; demanding better rather than more innovation.

Conclusions: deepening the Brussels effect?

The EU legal framework will shape the development, provision and deployment of RSs as it shapes global economic markets. As Bradford has convincingly argued, transnational corporations wishing to compete on the internal market of the EU will adapt their behaviour to stringent EU rules, because redesigning their internal processes per jurisdiction would be more costly [95]. It may be interesting, in that light, that corporate enterprise would prefer a level playing field that affords the development of responsible AI, noting that recently an 8.7 trillion USD investors alliance called on the EU legislature to integrate effective respect for digital rights in the DSA [96]. This may also help to address naïve or malicious invocations of cost-benefit analyses (e.g. [97]), which suggest that respect for human rights should be calculated and weighed against the costs of providing reliable and responsible RSs. This is not to suggest that compliance costs should not be assessed (see e.g. [98]), but to situate the role of such costs; imagine if we were to compare the costs of compliance with anti-corruption legislation against the benefits of corrupt governance.

This article, however, is not about cost-benefit analyses. It is about the assumptions that must be made when developing RSs, more specifically about the computational, machine-consumable proxies that determine the input, learning model and the output of RSs. In this article, I have explained how and why the issue of proxies is key to understanding both the productive nature and the potential corruption of RSs, due to their behaviourist and utilitarian underpinnings. Key concepts such as 'preferences' (whether latent or explicit) function as computable proxies for what individuals may desire or believe, compressing the rich and dynamic reality of human interaction to discrete sets of supposedly given (and manipulable) indicators of their future behaviours. Note that the dependencies this creates (on reductive understandings of the human condition) return in naïvely quantified cost-benefit analyses. Only if we become aware of the framing powers that are inherent in the choice and formalisation of the proxies, can the far-reaching implications of such choices be called out as both foundational and political. Legislation is needed to present those who call the shots on how to choose these proxies, with a well-thought-out choice architecture that constrains developers and providers in the right direction, aiming to protect human capabilities instead of reducing human agency.

As any computer scientist knows, it is constraints that afford freedom. Without constraints, a computing system cannot 'act'; computer code is all about constraints. What matters is which constraints, targeting what kind of behaviour, based on what types of feedback. In this article, I have argued that constraints do not only matter for computing systems, but also for economic markets that determine what type of products and what kind of providers are successful. In turn, legal constraints are the deep code of economic markets, which are not determined by the mythological narratives of Neo-Darwinism [74] but co-constituted by the legal constraints we decide to impose on them. Legal limits are key to abolishing the provision of

unreliable RSs that serve the preferences of advertiser intermediaries and/or the providers of very large online platforms instead of serving those in need of relevant and reliable recommendations.

References

- [1] L. Zou, L. Xia, Z. Ding, J. Song, W. Liu, en D. Yin, 'Reinforcement Learning to Optimize Long-term User Engagement in Recommender Systems', in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, New York, NY, USA, jul. 2019, pp. 2810–2818. doi: 10.1145/3292500.3330668.
- [2] S. Viljoen, J. Goldenfein, en L. McGuigan, 'Design choices: Mechanism design and platform capitalism', *Big Data & Society*, vol. 8, nr. 2, p. 20539517211034310, jul. 2021, doi: 10.1177/20539517211034312.
- [3] M. Hildebrandt, 'The issue of bias. The framing powers of machine learning', in *Machines We Trust: Perspectives on Dependable AI*, M. Pelillo en T. Scantamburlo, Red. Cambridge, Massachusetts: The MIT Press, 2021.
- [4] J. Dewey, 'Freedom and Culture', in *The Later Works of John Dewey, 1925 - 1953: 1938-1939, Experience and Education, Freedom and Culture, Theory of Valuation, and Essays*, SIU Press, 1988, p. 65.
- [5] M. Hildebrandt, 'Learning as a Machine. Crossovers Between Humans and Machines', *Journal of Learning Analytics*, vol. 4, nr. 1, pp. 6–23, mrt. 2017, doi: 10.18608/jla.2017.41.3.
- [6] N. Seaver, 'Captivating algorithms: Recommender systems as traps', *Journal of Material Culture*, vol. 24, nr. 4, pp. 421–436, dec. 2019, doi: 10.1177/1359183518820366.
- [7] F. Ricci, L. Rokach, en B. Shapira, Red., *Recommender Systems Handbook*, Softcover reprint of the original 2nd ed. 2015 edition. New York, NY: Springer, 2016.
- [8] D. Jannach, *Recommender Systems: An Introduction*, 1st edition. New York: Cambridge University Press, 2010.
- [9] A. Felfernig e.a., 'An overview of recommender systems in the internet of things', *J Intell Inf Syst*, vol. 52, nr. 2, pp. 285–309, apr. 2019, doi: 10.1007/s10844-018-0530-7.
- [10] L. Quijano-Sánchez, I. Cantador, M. E. Cortés-Cediel, en O. Gil, 'Recommender systems for smart cities', *Information Systems*, vol. 92, p. 101545, sep. 2020, doi: 10.1016/j.is.2020.101545.
- [11] S. Chadoulos, I. Koutsopoulos, en G. C. Polyzos, 'Mobile Apps Meet the Smart Energy Grid: A Survey on Consumer Engagement and Machine Learning Applications', *IEEE Access*, vol. 8, pp. 219632–219655, 2020, doi: 10.1109/ACCESS.2020.3042758.
- [12] J. T. Yun, C. M. Segijn, S. Pearson, E. C. Malthouse, J. A. Konstan, en V. Shankar, 'Challenges and Future Directions of Computational Advertising Measurement Systems', *Journal of Advertising*, vol. 49, nr. 4, pp. 446–458, aug. 2020, doi: 10.1080/00913367.2020.1795757.
- [13] M. Kranzberg, 'Technology and History: "Kranzberg's Laws"', *Technology and Culture*, vol. 27, pp. 544–560, 1986.
- [14] Y. Man, 'Towards a pluralistic epistemology: understanding the future of human technology interactions in shipping', *Doktorsavhandlingar vid Chalmers tekniska högskola. Ny serie*, nr. 4607, 2019, Geraadpleegd: 16 januari 2022. [Online]. Beschikbaar op: <https://trid.trb.org/view/1737757>
- [15] R. H. Thaler, C. R. Sunstein, en J. P. Balz, 'Choice Architecture', Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 1583509, apr. 2010. Geraadpleegd: 1 augustus 2016. [Online]. Beschikbaar op: <http://papers.ssrn.com/abstract=1583509>
- [16] M. Hildebrandt en S. Gutwirth, *Profiling the European Citizen. Cross-disciplinary Perspectives*. Dordrecht: Springer, 2008.
- [17] A. Vedder, 'KDD: The challenge to individualism', *Ethics and Information Technology*, vol. 1, pp. 275–281, 1999.

- [18] B. Custers, *The Power of Knowledge. Ethical, Legal, and Technological Aspects of Data Mining and Group Profiling in Epidemiology*. Nijmegen: Wolf Legal Publishers, 2004.
- [19] E. Bayamlioglu, I. Baraliuc, L. A. W. Janssens, en M. Hildebrandt, Red., *BEING PROFILED: COGITAS ERGO SUM: 10 Years of Profiling the European Citizen*. Amsterdam University Press, 2018.
- [20] F. Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and information*. Cambridge, Massachusetts London, England, 2016.
- [21] V. Eubanks, *Eubanks, V: Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York, NY, 2018.
- [22] S. U. Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York, 2018.
- [23] K. Crawford, *Atlas of AI: The Real Worlds of Artificial Intelligence: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven, 2021.
- [24] S. Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, 1st Edition. New York: PublicAffairs, 2019.
- [25] J. E. Cohen, *Between Truth and Power: The Legal Constructions of Informational Capitalism*. Oxford, New York: Oxford University Press, 2019.
- [26] K. Pistor, *The Code of Capital: How the Law Creates Wealth and Inequality*. Princeton: Princeton University Press, 2019.
- [27] M. Hildebrandt, *Law for Computer Scientists and Other Folk*. Oxford: Oxford University Press, 2020. Geraadpleegd: 10 augustus 2019. [Online]. Beschikbaar op: <https://global.oup.com/academic/product/law-for-computer-scientists-and-other-folk-9780198860884?cc=be&lang=en&>
- [28] G. E. M. Anscombe, 'On Brute Facts', *Analysis*, vol. 18, nr. 3, pp. 69–72, 1958, doi: 10.2307/3326788.
- [29] J. L. Austin, *How to Do Things with Words*, 2nd dr. Boston: Harvard University Press, 1975.
- [30] J. R. Searle, *Speech acts, an essay in the philosophy of language*. Cambridge, 1969.
- [31] N. MacCormick, *Institutions of Law: An EsSay in Legal Theory*, 1 édition. OUP Oxford, 2007.
- [32] M. Hildebrandt, 'The adaptive nature of text-driven law', *Journal of Cross-disciplinary Research in Computational Law*, sep. 2020, Geraadpleegd: 21 oktober 2020. [Online]. Beschikbaar op: <https://journalcrcl.org/crcl/article/view/2>
- [33] T. Saracevic, 'Relevance: A review of the literature and a framework for thinking on the notion in information science. Part II: nature and manifestations of relevance', *Journal of the American Society for Information Science and Technology*, vol. 58, nr. 13, pp. 1915–1933, 2007, doi: 10.1002/asi.20682.
- [34] T. J. Froehlich, 'Relevance reconsidered—Towards an agenda for the 21st century: Introduction to special topic issue on relevance research', *Journal of the American Society for Information Science*, vol. 45, nr. 3, pp. 124–134, 1994, doi: 10.1002/(SICI)1097-4571(199404)45:3<124::AID-ASI2>3.0.CO;2-8.
- [35] V. Packard, *The Hidden Persuaders*, Updated. Pocket, 1984.
- [36] A. Campagner, D. Ciucci, C.-M. Svensson, M. T. Figge, en F. Cabitza, 'Ground truthing from multi-rater labeling with three-way decision and possibility theory', *Information Sciences*, vol. 545, pp. 771–790, feb. 2021, doi: 10.1016/j.ins.2020.09.049.
- [37] G. H. Mead en C. W. Morris, *Mind, self, and society from the standpoint of a social behaviorist*. Chicago, Ill.: University of Chicago Press, 1962.
- [38] F. Stadler, 'From Methodenstreit to the "Science Wars" – an Overview on Methodological Disputes between the Natural, Social, and Cultural Sciences', *Formal and Informal Methods in Philosophy*, pp. 77–100, mrt. 2020, doi: 10.1163/9789004420502_006.
- [39] F. Maclachlan, 'Max Weber within the Methodenstreit', *Cambridge J Econ*, vol. 41, nr. 4, pp. 1161–1175, jul. 2017, doi: 10.1093/cje/bew051.
- [40] K. Hauser, 'Historical School and "Methodenstreit"', *Journal of Institutional and Theoretical Economics*, vol. 144, nr. 3, pp. 532–542, 1988.

- [41] J. Heath, 'Methodological Individualism', in *The Stanford Encyclopedia of Philosophy*, Summer 2020., E. N. Zalta, Red. Metaphysics Research Lab, Stanford University, 2020. Geraadpleegd: 11 september 2021. [Online]. Beschikbaar op: <https://plato.stanford.edu/archives/sum2020/entries/methodological-individualism/>
- [42] D. Ariely, *The Upside of Irrationality: The Unexpected Benefits of Defying Logic*, Reprint edition. New York etc.: Harper Perennial, 2011.
- [43] J. Frederik en M. Martijn, 'The new dot com bubble is here: it's called online advertising - The Correspondent', *The Correspondent*. Geraadpleegd: 11 september 2020. [Online]. Beschikbaar op: <https://thecorrespondent.com/100/the-new-dot-com-bubble-is-here-its-called-online-advertising/13228924500-22d5fd24>
- [44] M. Jesse en D. Jannach, 'Digital nudging with recommender systems: Survey and future directions', *Computers in Human Behavior Reports*, vol. 3, p. 100052, jan. 2021, doi: 10.1016/j.chbr.2020.100052.
- [45] C. R. Sunstein, *The Ethics of Influence: Government in the Age of Behavioral Science*. Cambridge University Press, 2016.
- [46] 'After GDPR, The New York Times cut off ad exchanges in Europe - and kept growing ad revenue', *Digiday*, 16 januari 2019. <https://digiday.com/media/gumgumtest-new-york-times-gdpr-cut-off-ad-exchanges-europe-ad-revenue/> (geraadpleegd 17 november 2019).
- [47] L. Johnson|March 1 en 2018, 'When Procter & Gamble Cut \$200 Million in Digital Ad Spend, It Increased Its Reach 10%'. <https://www.adweek.com/brand-marketing/when-procter-gamble-cut-200-million-in-digital-ad-spend-its-marketing-became-10-more-effective/> (geraadpleegd 17 november 2019).
- [48] N. Lomas, 'The case against behavioral advertising is stacking up | TechCrunch', *TechCrunch*, 20 januari 2019. https://techcrunch.com/2019/01/20/dont-be-creepy/?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlLnNvbS8&guce_referrer_sig=AQAAAL9IDyWhilwkpISW_ZaHtaZ6CH2cvlxxatXP7wGckRUOankBY6_WbkkHxwok20WAljppLKMAL37TrruSV58k6iJSIJQ1O2BAMdgeo8bn0angcirt6VL8dTVdeE4pE6BJn6yoiYRVjCWyAKy3vETkTA7cMfgeJCg6s7P9u3ZUek7 (geraadpleegd 11 september 2020).
- [49] M. 29th 2019 10:44am-M. M. Fri, 'What If Google And Facebook Admitted That All This Ad Targeting Really Doesn't Work That Well?', *Techdirt*. <https://www.techdirt.com/articles/20190319/23223441830/what-if-google-facebook-admitted-that-all-this-ad-targeting-really-doesnt-work-that-well.shtml> (geraadpleegd 11 september 2021).
- [50] V. Marotta, V. Abhishek, en A. Acquisti, 'Online Tracking and Publishers' Revenues: An Empirical Analysis'.
- [51] 'Norway: Norwegian Consumer Council calls for ban of surveillance-based advertising', *DataGuidance*, 6 juli 2021. <https://www.dataguidance.com/news/norway-norwegian-consumer-council-calls-ban> (geraadpleegd 11 september 2021).
- [52] 'Facebook Disables Accounts Tied to NYU Research Project', *Bloomberg.com*, 5 augustus 2021. Geraadpleegd: 16 september 2021. [Online]. Beschikbaar op: <https://www.bloomberg.com/news/articles/2021-08-03/facebook-disables-accounts-tied-to-nyu-research-project>
- [53] N. Watzman, 'The political ads Facebook won't show you', *Cybersecurity for Democracy*, 12 mei 2021. <https://medium.com/cybersecurity-for-democracy/the-political-ads-facebook-wont-show-you-e0d6181bca25> (geraadpleegd 16 september 2021).
- [54] C. Faife, 'Facebook Rolls Out News Feed Change That Blocks Watchdogs from Gathering Data – The Markup', *The Markup*. <https://themarkup.org/citizen-browser/2021/09/21/facebook-rolls-out-news-feed-change-that-blocks-watchdogs-from-gathering-data> (geraadpleegd 21 september 2021).
- [55] M. Hildebrandt, 'Behavioural Biometric Profiling and Transparency Enhancing Tools', *Future of Identity in Information Society*, Brussel, 2009. [Online]. Beschikbaar op: www.fidis.net

- [56] M. Hildebrandt, *Smart Technologies and the End(s) of Law. Novel Entanglements of Law and Technology*. Cheltenham: Edward Elgar, 2015.
- [57] A. Popescu e.a., 'Increasing Transparency and Privacy for Online Social Network Users – USEMP Value Model, Scoring Framework and Legal', in *Privacy Technologies and Policy: Third Annual Privacy Forum, APF 2015, Luxembourg, Luxembourg, October 7-8, 2015, Revised Selected Papers*, B. Berendt, T. Engel, D. Ikonomou, D. Le Métayer, en S. Schiffner, Red. Cham: Springer International Publishing, 2016, pp. 38–59. doi: 10.1007/978-3-319-31456-3_3.
- [58] H. G. Frankfurt, *On Bullshit*, 1 edition. Princeton, NJ: Princeton University Press, 2005.
- [59] S. Vaidhyanathan, *Antisocial Media: How Facebook Disconnects Us and Undermines Democracy*. New York, NY, United States of America: Oxford University Press, 2018.
- [60] K. Roose, 'Inside Facebook's Data Wars', *The New York Times*, 14 juli 2021. Geraadpleegd: 16 januari 2022. [Online]. Beschikbaar op: <https://www.nytimes.com/2021/07/14/technology/facebook-data.html>
- [61] M. Callon en L. J., 'On qualculation, agency, and otherness', *Environment and Planning D: Society and Space*, vol. 23, nr. 5, pp. 717–733, 2005.
- [62] M. Polanyi, *The Tacit Dimension*. Garden City, New York: Anchor Books, 1966.
- [63] W. B. Gallie, 'Essentially Contested Concepts', *Proc. Aristotelian Soc'ty*, vol. 56, pp. 167–198, 1956.
- [64] P. Jones, 'Refugees help power machine learning advances at Microsoft, Facebook, and Amazon', *Rest of World*, 22 september 2021. <https://restofworld.org/2021/refugees-machine-learning-big-tech/> (geraadpleegd 30 september 2021).
- [65] E. Denton, A. Hanna, R. Amironesei, A. Smart, en H. Nicole, 'On the genealogy of machine learning datasets: A critical history of ImageNet', *Big Data & Society*, vol. 8, nr. 2, p. 20539517211035956, jul. 2021, doi: 10.1177/20539517211035955.
- [66] C. Goodhart, 'Problems of monetary management: the U.K. experience', *Papers in monetary economics 1975 ; 1*, vol. 1, 1975.
- [67] D. T. Campbell, 'Assessing the impact of planned social change', *Evaluation and Program Planning*, vol. 2, nr. 1, pp. 67–90, jan. 1979, doi: 10.1016/0149-7189(79)90048-X.
- [68] M. Strathern, "'Improving ratings": audit in the British University system', *European Review*, vol. 5, nr. 3, pp. 305–321, jul. 1997, doi: 10.1002/(SICI)1234-981X(199707)5:3<305::AID-EURO184>3.0.CO;2-4.
- [69] J. Stray, 'Designing Recommender Systems to Depolarize', *CoRR*, vol. abs/2107.04953, 2021, Geraadpleegd: 16 januari 2022. [Online]. Beschikbaar op: <https://arxiv.org/abs/2107.04953>
- [70] D. Kahneman, 'Maps of Bounded Rationality: Psychology for Behavioral Economics', *The American Economic Review*, vol. 93, nr. 5, pp. 1449–1475, 2003.
- [71] P. M. Todd en G. Gigerenzer, 'Environments That Make Us Smart: Ecological Rationality', *Curr Dir Psychol Sci*, vol. 16, nr. 3, pp. 167–171, jun. 2007, doi: 10.1111/j.1467-8721.2007.00497.x.
- [72] J. Favini, 'What if Competition Isn't As "Natural" As We Think?', *Slate*, 23 januari 2020. Geraadpleegd: 19 september 2021. [Online]. Beschikbaar op: <https://slate.com/technology/2020/01/darwin-competition-collaboration-evolutionary-biology-climate-change.html>
- [73] M. Sahlins, *Stone age economics*. London: Tavistock, 1974.
- [74] M. D. Sahlins, *The Use and Abuse of Biology: An Anthropological Critique of Sociobiology*, 1st edition. Ann Arbor, Mich: University of Michigan Press, 1976.
- [75] T. Mitchell, *Machine Learning*, 1 edition. New York: McGraw-Hill Education, 1997.
- [76] R. Brooks, 'Machine Learning Explained', *Robots, AI, and other stuff*, 28 augustus 2017. <http://rodnebrooks.com/forai-machine-learning-explained/> (geraadpleegd 29 augustus 2017).
- [77] J. M. Hofman, A. Sharma, en D. J. Watts, 'Prediction and explanation in social systems', *Science*, vol. 355, nr. 6324, pp. 486–488, feb. 2017, doi: 10.1126/science.aal3856.

- [78] Z. C. Lipton en J. Steinhardt, 'Research for practice: troubling trends in machine-learning scholarship', *Commun. ACM*, vol. 62, nr. 6, pp. 45–53, mei 2019, doi: 10.1145/3316774.
- [79] J. Dewey, *The public & its problems*. Chicago: The Swallow Press, 1927.
- [80] J. Gibson, *The Ecological Approach to Visual Perception*. New Jersey: Lawrence Erlbaum Associates, 1986.
- [81] M. Heras-Escribano, *The Philosophy of Affordances*, 1st ed. 2019 edition. Palgrave Macmillan, 2019.
- [82] D. A. Norman, 'Affordance, Conventions, and Design', *interactions*, vol. 6, nr. 3, pp. 38–43, mei 1999, doi: 10.1145/301153.301168.
- [83] A. Sen, *Commodities and Capabilities*. Oxford University Press, USA, 1999.
- [84] I. Robeyns, 'The Capability approach: a theoretical survey', *Journal of Human Development*, vol. 6, nr. 1, pp. 93–114, 2005.
- [85] M. Hildebrandt en L. Tielemans, 'Data protection by design and technology neutral law', *Computer Law & Security Review*, vol. 29, nr. 5, pp. 509–521, okt. 2013, doi: 10.1016/j.clsr.2013.07.004.
- [86] A. Bradford, 'The Brussels Effect', *Northwestern University Law Review*, vol. 107, nr. 1, pp. 1–68, 2012.
- [87] L. A. Bygrave, 'The "Strasbourg Effect" on data protection in light of the "Brussels Effect": Logic, mechanics and prospects', *Computer Law & Security Review*, vol. 40, p. 105460, apr. 2021, doi: 10.1016/j.clsr.2020.105460.
- [88] E. A. Brasher, 'Addressing the Failure of Anonymization: Guidance from the European Union's General Data Protection Regulation', *Colum. Bus. L. Rev.*, vol. 2018, p. 209, 2018, [Online]. Beschikbaar op: <https://heinonline.org/HOL/Page?handle=hein.journals/colb2018&id=215&div=&collection=>
- [89] M. C. Schraefel, R. Gomer, E. Gerding, en C. Maple, 'Rethinking transparency for the Internet of Things', *Life and the Law in the Era of Data-Driven Agency*, jan. 2020, Geraadpleegd: 11 september 2020. [Online]. Beschikbaar op: <https://www.elgaronline.com/view/edcoll/9781788971997/9781788971997.00012.xml>
- [90] P. Ohm, 'Throttling machine learning', *Life and the Law in the Era of Data-Driven Agency*, jan. 2020, Geraadpleegd: 11 september 2020. [Online]. Beschikbaar op: <https://www.elgaronline.com/view/edcoll/9781788971997/9781788971997.00019.xml>
- [91] L. Edwards en M. Veale, 'Slave to the Algorithm? Why a "Right to Explanation" is Probably Not the Remedy You are Looking for', Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 2972855, mei 2017. Geraadpleegd: 14 juni 2017. [Online]. Beschikbaar op: <https://papers.ssrn.com/abstract=2972855>
- [92] F. Xu, H. Uszkoreit, Y. Du, W. Fan, D. Zhao, en J. Zhu, 'Explainable AI: A Brief Survey on History, Research Areas, Approaches and Challenges', in *Natural Language Processing and Chinese Computing*, Cham, 2019, pp. 563–574. doi: 10.1007/978-3-030-32236-6_51.
- [93] R. Binns en M. Veale, 'Is That Your Final Decision? Multi-Stage Profiling, Selective Effects, and Article 22 of the GDPR'. SocArXiv, 10 augustus 2021. doi: 10.31235/osf.io/7mq6z.
- [94] M. Hildebrandt, 'Feedback on the proposed Artificial Intelligence Act by Mireille Hildebrandt, Vrije Universiteit Brussel'. 19 juli 2021. Geraadpleegd: 18 januari 2022. [Online]. Beschikbaar op: https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12527-Artificial-intelligence-ethical-and-legal-requirements/F2662611_en
- [95] A. Bradford, *The Brussels Effect: How the European Union Rules the World*. New York: Oxford University Press, 2020.
- [96] Investor Alliance for Human Rights, 'Investors Representing over US\$8.7T Call on the European Union to Develop Digital Rights Regulation'. <https://investorsforhumanrights.org/news/investors-representing-over-us87t-call-european-union-develop-digital-rights-regulation> (geraadpleegd 18 januari 2022).

- [97] M. Laurer, A. Renda, en T. Yeung, 'Clarifying the costs for the EU's AI Act', *CEPS*, 24 september 2021. <https://www.ceps.eu/clarifying-the-costs-for-the-eus-ai-act/> (geraadpleegd 18 januari 2022).
- [98] CEPS e.a., *Study to support an impact assessment of regulatory requirements for Artificial Intelligence in Europe: final report*. LU: Publications Office of the European Union, 2021. Geraadpleegd: 18 januari 2022. [Online]. Beschikbaar op: <https://data.europa.eu/doi/10.2759/523404>