



“Please understand we cannot provide further information”: evaluating content and transparency of GDPR-mandated AI disclosures

Alexander J. Wulf¹ · Ognyan Seizov¹

Received: 13 September 2021 / Accepted: 25 February 2022 / Published online: 11 May 2022
© The Author(s) 2022

Abstract

The General Data Protection Regulation (GDPR) of the EU confirms the protection of personal data as a fundamental human right and affords data subjects more control over the way their personal information is processed, shared, and analyzed. However, where data are processed by artificial intelligence (AI) algorithms, asserting control and providing adequate explanations is a challenge. Due to massive increases in computing power and big data processing, modern AI algorithms are too complex and opaque to be understood by most data subjects. Articles 15 and 22 of the GDPR provide a modest regulatory framework for automated data processing by, among other things, mandating that data controllers inform data subjects about when it is being used, and its logic and ramifications. Nevertheless, due to the phrasing of the articles and the numerous exceptions they allow, doubts have arisen about their effectiveness. In this paper, we empirically evaluate the quality and effectiveness of AI disclosures as mandated by the GDPR. By means of an online survey ($N=835$), we investigated how data subjects expect to be informed about the automated processing of their data. We then conducted a content analysis of the AI disclosures of $N=100$ companies and organizations. The combined findings reveal that current GDPR-mandated disclosures do not meet the expectations and needs of data subjects. Explanations drawn up following the guidelines of the generic formulations of the GDPR differ widely and are often vague, incomplete and lack transparency. In our conclusions we identify a path towards standardizing and optimizing AI information notices.

Keywords Artificial intelligence · Automated data processing · Empirical legal studies · Explainable AI · GDPR · Information disclosures

1 Introduction

Artificial intelligence (AI) algorithms find application in a variety of business fields.¹ The exponential growth in computing power, the widespread adoption of cloud computing, and the presence of ever more complex algorithms hidden behind user-friendly interfaces² have cemented the place of

AI in the centre of the datafied³ digital economy. The flipside of this ‘algorithmic boom’ is the diminishing explicability of the inner workings of automated systems: while classical analytical and predictive models, such as simple regressions and correlations, can easily be explained and reconstructed by a human being, modern approaches, such as neural networks or random forests, operate in what is often called a ‘black box’,

¹ Faust and Schäfer (2019).

² Buyers (2018), Finlay (2017).

³ Van Dijck (2014).

✉ Alexander J. Wulf
alexander.wulf@srh.de

¹ Berlin School of Management, SRH Berlin University of Applied Sciences, Ernst-Reuter-Platz 10, 10587 Berlin, Germany

in which data processing goes through too many steps and loops to be comprehensible to human beings.⁴ At the same time, computer scientists, sociologists, ethicists and also various international organizations have repeatedly called for the development and implementation of explainable AI that promotes equality and prosperity.⁵ The General Data Protection Regulation (GDPR)⁶ of the European Union is among the first examples of legislation that addresses the problem of AI explicability and ethics. As thoroughly reviewed elsewhere, Articles 15 and 22 of the GDPR, in particular, act as a modest regulatory framework for automated data processing and decision-making aimed at natural persons (i.e., data subjects).⁷ The goal of the regulation in this instance is to reconcile the extensive application of new technologies and big data⁸ with the fundamental human right to the protection of personal data.⁹ As Wulf and Seizov point out in their review, however, the requirements set out in Articles 15 and 22 most likely do not go far enough to afford data subjects sufficient control over the automated processing of their personal information.¹⁰ In this paper, we report on an empirical legal study¹¹ in which we took a first step towards testing the effectiveness and limits of the GDPR in that regard.

The time is ripe to address this problem. The GDPR has been in force since May 2018;¹² the use of algorithms has been growing steadily in a variety of consumer contexts;¹³ international researchers in the fields of computer science,¹⁴ ethics,¹⁵ law¹⁶ and empirical legal studies¹⁷ have called for more clarity and more concrete oversight of automated data processing. Furthermore, recent case law from the Netherlands demonstrates the growing practical relevance that explainable AI has for the various stakeholders of

technology companies.¹⁸ In our study of the effectiveness of GDPR-mandated disclosures on the automated processing of personal data, we addressed four research questions which we elaborate below.

Two of the defining features of the algorithmic boom are its extensive application across a variety of industries and the towering complexity of its calculations.¹⁹ Both of these characteristics pose practical challenges to consumer understanding. Previous experimental research has shown that even when asked to read less technical disclosures such as terms and conditions or privacy notices, consumers tend to misunderstand or ignore them.²⁰ Marotta-Wurgler confirmed such experimental accounts by an analysis of 48,000 online shoppers' clickstream data (i.e., their online shopping pathway, click by click) which showed that they spent negligibly short amounts of time on the information disclosure webpages.²¹ In a qualitative study on transparent online disclosure in the EU, Wulf and Seizov found that consumers habitually click away cookie notices because they find them too tedious²²— so what hope is left for a much more complex explanation of AI applications for data processing and automated decision-making? Our first guiding question is therefore: *How do consumers perceive AI usage and how do they wish to be informed about it?*

Consumers' expectations of AI disclosure are naturally related to the actual AI usage in which companies engage. The literature on AI deployment in society paints a picture of omnipresence and opacity. Buyers, Cooper, and Finlay describe a multitude of algorithms that carry out a growing number of tasks and feed on a seemingly endless stream of data generated by billions of connected devices.²³ Helbling and colleagues question whether democracy can survive the massive deployment of algorithms that collect and analyze massive amounts of private information.²⁴ Van Boom and colleagues focus on algorithmic price discrimination. Informing consumers diligently about its negative outcomes produces strong negative reactions, pointing towards a potential disincentive for businesses to disclose diligently.²⁵ In fact, many AI-powered businesses skirt transparent disclosure and employ 'dark patterns' of explicability.²⁶ Instead of disclosing what algorithms are employed, how they work and for what purpose, many businesses intentionally obscure that information through misleading formulations, loaded

⁴ Adadi and Berrada (2018), Buyers (n 2); Cooper (2018). Regarding the impossibility for users to trace the (AI-assisted) decisions of media intermediaries about the presentation of online content, see Schwartzmann et al. (2020).

⁵ Bhatt (2018), Floridi et al. (2018), High-Level Expert Group on AI (2019), Olhede and Wolfe (2018), Shahriari and Shahriari (2017), Suzuki (2018), Whittlestone et al. (2019)

⁶ EU General Data Protection Regulation (GDPR) (2016).

⁷ Wulf and Seizov (2020a, b).

⁸ Recital 6 (GDPR).

⁹ Recital 1 (GDPR). Besides the 'protection of personal data', we subsequently at times also refer to 'data privacy', as does much of the corporate communication we quote. While some consider these two concepts to be synonymous, 'data privacy' is arguably a broader term, encompassing also the protection of non-personal data.

¹⁰ Wulf and Seizov (n 7).

¹¹ Wulf (2016).

¹² Rossow (2018).

¹³ Faust and Schäfer (n 1).

¹⁴ Adadi and Berrada (n 4); Bhatt (n 5).

¹⁵ Floridi et al. (n 5).

¹⁶ Feiler et al. (2018).

¹⁷ Wulf and Seizov (n 7).

¹⁸ Gellert et al. (2021).

¹⁹ Buyers (n 2); Cooper (n 4); Finlay (n 2).

²⁰ Ben-Shahar and Chilton (2016).

²¹ Marotta-Wurgler (2011).

²² Wulf and Seizov (2020a, b).

²³ Buyers (n 2); Cooper (n 4); Finlay (n 2).

²⁴ Helbling et al. (2019).

²⁵ van Boom et al. (2020).

²⁶ Chromik et al. (2019).

user dialogues, and false choices that create the illusion of information and consumer control. Given this atmosphere of lack of clarity and doubt surrounding the application of AI for automated data processing and decision-making, we next tackled the question: *How is AI currently used by companies operating in Germany?*

The combination of consumers' tendency to bypass legal information online and businesses' imperfect records of transparent disclosure makes the effectiveness of the GDPR-mandated right to AI disclosure and explanation a point of concern. Wulf and Seizov reviewed a number of cases related to Articles 15 and 22 GDPR or to their precursors in German national law and identified a number of instances in which German, Dutch and other European courts limited data subjects' rights of access to information, be it to reduce administrative burden, to keep companies' proprietary algorithms and internal processes hidden (i.e., trade secrets), or to narrow the very definition of what constitutes 'personal data'.²⁷

Three recent judgements passed by the Amsterdam District Court of first instance in March 2021 paint a similarly mixed picture of the degrees of restrictiveness of Dutch courts when it comes to granting data subjects access to information about the algorithms that affect them. In two connected cases, the court denied data subjects access to this information, whereas in another case it granted it. In a judgement involving the ride-hailing company Uber, some of the company's drivers asked for meaningful information about the logic involved in the algorithm used to match drivers and passengers. The court followed Uber's argument that the algorithm in question does not have any legal consequences for the driver, nor does it significantly affect them in any other way. The court, therefore, ruled that no automated decision-making within the meaning of Article 22 GDPR takes place and thus rejected the data subjects' request for information about this algorithm.²⁸ In a second, connected judgement by the Amsterdam District Court which also involved Uber, some of its drivers asked for information about the logic involved in an algorithmic decision that led to the cancellation of their contracts due to fraudulent behaviour. The company's privacy statement indicated that such decisions are taken fully automatically. Uber argued that contrary to this statement, in the EU and in the UK such decisions are not solely based on automated decision-making. According to the court Uber was able to convincingly make the case (which was also not disputed by the plaintiffs) that the decision to cancel the drivers' contracts involved significant human intervention. The court,

therefore, denied the data subjects' requests for access to further information about the algorithms in this case also.²⁹

The fact that the phrasing of Articles 15 and 22 GDPR allows companies to successfully avoid making thorough information disclosures led Gierschmann and colleagues to characterize them as "an incomplete norm".³⁰ Feiler and colleagues also characterize the GDPR standard for compliance in this instance as vague, particularly in comparison with previous consumer-information legislation, such as the now repealed Data Protection Directive of the EU that came into force in 1995.³¹ Taking all the above concerns into account, Wachter and colleagues felt compelled to point out that the GDPR does not actually establish a right to explanation of algorithmic decision-making.³²

However, in a third judgement passed in March 2021, the Amsterdam District Court recognized that the GDPR establishes such a right, though only if certain strict conditions are met. It argued that the rights of the drivers of the ride-hailing company Ola were significantly negatively affected by a fully automated algorithm that fined drivers without human intervention for invalid rides. The court argued that the provisions of the GDPR prohibit the company from subjecting data subjects to the consequences of such algorithmic decision-making and that this would only be permissible if it were necessary for the performance of the contract or the company had obtained their explicit prior consent. Because the company was unable to show that these exceptions applied, the court ordered it to provide information about the logic involved in this algorithm.³³

The effectiveness of Articles 15 and 22 GDPR is in any case diminished by the fact that data subjects need to take action to inform themselves about their rights and then to also make use of them. Wulf and Seizov have called this effect "a double transparency barrier".³⁴ In this context, we sought to address the question: *Is the right to AI disclosure laid down in the GDPR effective in informing consumers?*

The rule that information disclosures must be clear and understandable has long been enshrined in EU law, albeit in the most general terms.³⁵ The practical aspects of transparency have habitually been under-regulated and under-defined, and the GDPR seems to continue this trend.³⁶ Previous research into the understandability of information

²⁷ Wulf and Seizov (n 7).

²⁸ Uber drivers v. Uber B.V. C/13/687315 / HA RK 20–207, District Court, Amsterdam (11–03-2021).

²⁹ Uber drivers v. Uber B.V. C/13/692003 / HA RK 20–302, District Court, Amsterdam (11–03-2021).

³⁰ Gierschmann et al. (2017).

³¹ Feiler et al. (n 16).

³² Wachter (2017).

³³ Ola drivers v. Ola Netherlands B.V. C/13/689705 / HA RK 20–258, District Court, Amsterdam (11–03-2021).

³⁴ Wulf and Seizov (n 7).

³⁵ Luzak (2014), Seizov et al. (2019), Wulf (2014)

³⁶ Feiler et al. (n 16).

disclosures has operationalized the concept in different ways. Furnell and Phippen focused on vocabulary and evaluating information notices according to reading-level as defined by the Fleisch–Kincaid scale. They found several terms and conditions of social media companies to be very hard to comprehend.³⁷ Pollach examined the linguistic and syntactical aspects of nearly two dozen online shops and uncovered numerous instances of obfuscation due to uncertain modal phrasing, convoluted sentence structures and imprecise use of adverbs.³⁸ Waller emphasized the importance of layout and noted how many electronic texts lack the clear structure and typography that are normally an essential feature of printed documents.³⁹ Last but not least, visuals can help communicate complex concepts more clearly thanks to multimodal meaning-making, if used correctly and adeptly.⁴⁰ While legal documents are traditionally text-only, innovative suggestions such as the one put forward by Berger-Walliser et al. provide a roadmap for the implementation of visual elements into legal document design.⁴¹ Studies by Seizov and Wulf have shown that many businesses are eager to experiment with alternative disclosure formats which have the potential to make complex legal and technical information more accessible.⁴² However, in the absence of legislation that explicitly allows the use of visualization techniques in the provision of consumer information, breaking away from the text-only norm creates legal risks which ultimately disincentivize innovation in this field. Like previous EU legislation, the GDPR, therefore, remains vague on what exactly constitutes clear and understandable disclosures. We thus rely on the multidisciplinary criteria set out by Seizov and colleagues⁴³ to answer the final guiding question: *Are the AI disclosures mandated by the GDPR transparent?*

The rest of the paper is dedicated to addressing these four guiding questions. In our study we used a combination of methods and data sources to investigate them. To answer the first question, we analyzed the opinions of 835 online survey respondents in order to establish a baseline of experience with and expectations of automated processing, its applications and explanations of it (see Sect. 2). To address the remaining questions, we exercised the right to access information granted to us by the GDPR and collected AI disclosures from 100 companies and organizations which we content-analyzed (see Sect. 3). In Sect. 4, we provide a summary evaluation of the consumers' expectations and how the companies operationalized the data processing disclosures

as required by the GDPR. Finally, in Sect. 5 we suggest how information requirements in this area could be rendered more concrete and practical.

2 A survey of consumers' knowledge and experience of AI in e-commerce

2.1 Method

We begin with an overview of consumers' knowledge and experience in matters of AI-automated data processing and the right to transparent information, which the European Commission included in the provisions of the GDPR. We also took the opportunity to collect the respondents' opinions on how they would like to be informed about the automated processing of their data, a perspective that has not been taken into account anywhere near often enough, either by legislators or by disclosers. In a sub-section of a large-scale online survey on the transparency of online contracts we explored the respondents' awareness of AI usage and their experience and knowledge of the GDPR as related to their rights to data access and to explanation of decisions made by AI algorithms. In what follows we present the data from the AI section of this survey only. We surveyed a sample of 835 British respondents whom we gathered through a participant recruitment company that specializes in academic studies (average age in years $M=36.72$, $SD=10.68$, female—66.2 percent, university-educated—57 percent, shop for goods and services online 'often' or 'very often'—77.3 percent). We consider our study exploratory, because the recruitment company's sampling strategy did not guarantee that our sample exactly matches the demographics of the average UK or EU online shopper. However, according to recent data published by Eurostat, the European Statistical Office, our sample is close to being representative of EU online shoppers, who are on average young and highly educated, as in our sample.⁴⁴ As a brief introduction to the topic for the survey participants, we pointed out that “[o]nline businesses use artificial intelligence (AI) algorithms for various purposes” and then asked the participants to “[p]lease tell us a little bit about your experience and opinion regarding this topic.”

2.2 Results

To establish a baseline of AI exposure, we began by asking the respondents questions about their experience with AI in their daily lives (see Table 1). Over half of them did not know whether they had used an AI-powered service and another 18 percent denied ever having used one. Similarly,

³⁷ Furnell and Phippen (2012).

³⁸ Pollach (2005).

³⁹ Waller (2017).

⁴⁰ Bateman et al. (2017), Seizov and Wildfeuer (2017).

⁴¹ Berger-Walliser et al. (2017)

⁴² Seizov and Wulf (2020).

⁴³ Seizov et al. (n 35); Seizov and Wulf (n 42).

⁴⁴ Eurostat (2021).

Table 1 Experiences with and attitudes towards automated decision-making

	Yes (%)	No (%)	Don't know (%)
Have you ever used a service that was powered by an Artificial Intelligence (AI) algorithm?	30.8	18.4	50.8
Has an AI algorithm ever made an autonomous decision about you?	21.6	22.2	56.3
Should businesses be allowed to use AI algorithms to make decisions about you without informing you about it?	17.2	64.8	18.0
Do you believe an AI algorithm makes more objective decisions than a human?	34.5	37.5	28.5

N = 835 respondents

56 percent of the respondents did not know whether they had been subject to automated decision-making, with the rest being more or less evenly split. The large amounts of ‘don’t know’ answers indicate that artificial intelligence, though abundantly used and increasingly crucial in many daily activities and commercial transactions, has not yet become a central issue in public awareness. They also demonstrate that the status quo falls short of the AI transparency standard put forward by the EU’s High Level Expert Group on AI, which includes the requirement that consumers be explicitly notified of all instances where they will encounter and interact with an algorithm.⁴⁵

In contrast to the level of uncertainty regarding exposure to AI, the majority of the respondents were clearly opposed to being subjected to automated decision-making without prior information—nearly 65 percent stated that view, against 17 percent who would not mind and 18 percent who did not have a firm opinion. Regarding the question whether the respondents thought that AI algorithms make more objective decisions than humans, the responses were rather evenly split. These results reveal something of a knowledge gap regarding AI applications, but they also show that consumers would, on average, like to be informed about the automated processes that affect them.

We then asked the respondents whether decisions made by an algorithm should be reviewed by a human and, if yes, under what circumstances.⁴⁶ (We report on these findings in text-only form in this paragraph.) Only four percent reported that they would accept automated decision-making without human review. A total of 21 percent said human review was necessary if a consumer actively complains about a decision, and 30 percent said any negative AI decision should automatically trigger human review. The remaining 45 percent would require human review of *all* automated decisions, regardless of their outcome. Since the right to

explanation and human review of algorithmic decision-making is an important part of the GDPR, we then inquired about the respondents’ knowledge of and experience with this legal statute.⁴⁷ Thirteen percent of the respondents did not know they had the right to request information about how their personal data are being processed and 77 percent had never felt the need to make such a request. Thus, nine-tenths of our sample had not made use of the generous data subject rights the GDPR affords them. Of the remaining 10 percent, seven had made at least one subject access request, with or without specific questions regarding AI usage, and were satisfied with the results, while three percent did not find the response that they received helpful. Thus, our findings reveal another knowledge gap, namely the divide between those who believed in the need for human review over AI decisions (96 per cent) vs. those who did not know they could or did not wish to exercise their right to request an explanation about how their personal data is being processed, by humans and by algorithms alike (90 percent).

Finally, we asked the respondents to indicate how detailed the AI disclosures should be. The results (see Fig. 1) speak for a rather detailed disclosure that alerts consumers to the fact that they are being subjected to an automated decision, explains the basic logic the algorithm employs and lists the personal data that flow into the automated decision-making process (44 percent). Fifteen percent would go even further and would like to see the computer code that powers the algorithm, something the GDPR does not currently foresee. The remaining respondents would be satisfied with a simple disclosure that informs them that an algorithm is at work, with (20 percent) or without (21 percent) an explanation of its underlying logic.

⁴⁵ High-Level Expert Group on AI (n 5).

⁴⁶ The question used to gather this data was: “Should decisions made by an Artificial Intelligence (AI) algorithm be reviewed by a human?”, with the response options “No.”, “Yes, but only if a consumer complains about the decision.”, “Yes, but only if a consumer is negatively affected by the decision.” and “Yes, always.”

⁴⁷ The question used to gather this data was: “Since May 2018, the General Data Protection Regulation (GDPR) grants you the right to ask companies to inform you about the amount, kind, and usage of the personal data they hold about you. Have you ever taken advantage of that right?”, with the response options “No – because I did not know that I had this right at all.”, “No – because I haven’t felt the need to take advantage of this right.”, “Yes – but I was not satisfied with the information I received.” and “Yes – and I was satisfied with the information I received.”

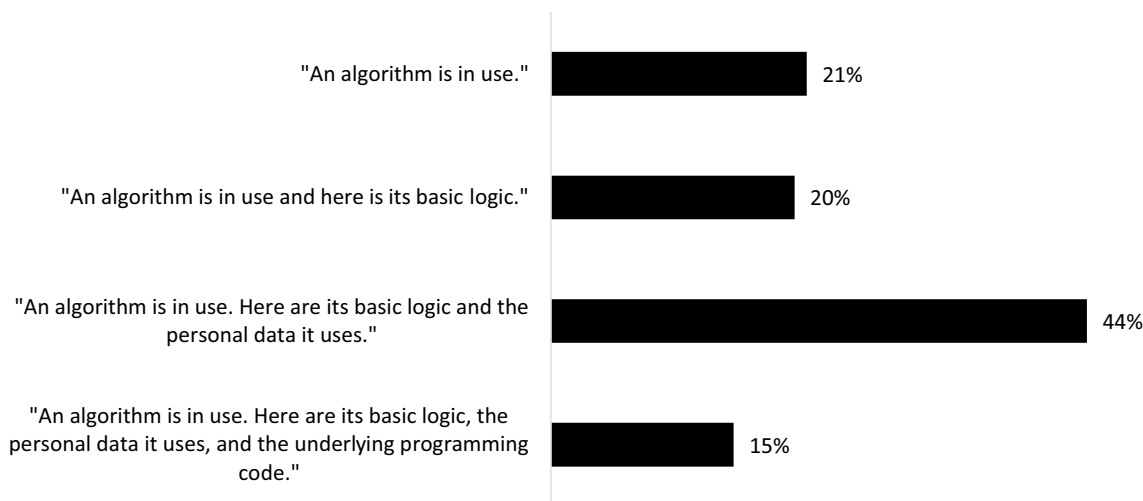


Fig. 1 Levels of detail that the respondents expect of an AI disclosure. $N=835$ respondents. The question used to gather this data was: “How much do you want to know about the Artificial Intelligence (AI) algorithms businesses employ to make decisions about you?”

Thus almost 80 percent of the respondents required a substantive AI disclosure and more than half of them expected information not only on the logic but also on the personal data that go into the decision-making process. This finding echoes the recommendation of the IEEE (Institute of Electrical and Electronics Engineers) that the basis of every algorithmic decision be clearly identifiable.⁴⁸ It also confirms the viability of a recent proposal made by Wulf and Seizov for an ‘AI Explainability One-Pager’ that reveals the algorithm’s basic logic, its input data and the way in which it assigns individual users to different groups.⁴⁹ With these insights on consumer experiences with, and expectations of, the AI disclosure process, we proceed to the second part of our study.

3 Exercising the right to information and explanation: a content analysis of 100 AI disclosures

3.1 Method

In the following section, we present our sample of 100 companies and organizations. With each of them, we exercised our right as per the GDPR to information on the nature, extent and consequences of AI, algorithms and automated decisions about us as consumers of their products or users of their services. On the one hand this is a convenience sample⁵⁰, given that we sampled companies and organizations

with which our research team and colleagues already had contractual relationships, i.e., accounts, subscriptions, past purchases and the like. On the other hand the list includes 11 of the Top 100 German companies according to revenues⁵¹; it also features 10 of the Top 100 US companies according to Forbes⁵² and four major players with non-German EU origins.⁵³ The sample counterbalances these ‘heavy-weights’ with smaller companies with various degrees of market power and capitalization on the German, European and global markets; it also features various branches of the local German administrations such as tax authorities, police, regulatory and judicial authorities. Figure 2 presents the full distribution of the organizations across different industries. While the sample thus reflects the range of companies and organizations offering goods and services online to consumers in Germany and we took care to sample within industries with highly standardized products and similar data processing practices (e.g., banking, insurance, air travel, e-commerce, etc.), we cannot claim that the sample is fully representative, so any findings derived from it cannot be generalised beyond the specific set of companies that we investigate and the results should thus be regarded as exploratory.⁵⁴

Table 2 provides further sample characteristics. German corporations and foreign legal forms formed the bulk of the

⁴⁸ Shahriari and Shahriari (n 5).

⁴⁹ Wulf and Seizov (2020a).

⁵⁰ Etikan et al. (2016).

⁵¹ Statista (2020a, b).

⁵² Forbes (2020).

⁵³ Statista (2020a, b).

⁵⁴ To obtain a fully representative sample, e.g., for the EU economy, the sampling would have to be based on an official statistical classification scheme such as Eurostat’s NACE Rev. 2. However, this would by far exceed the scope of this study.

Fig. 2 Industries represented in the sample. *N*= 100 companies and organizations (exploratory sample)

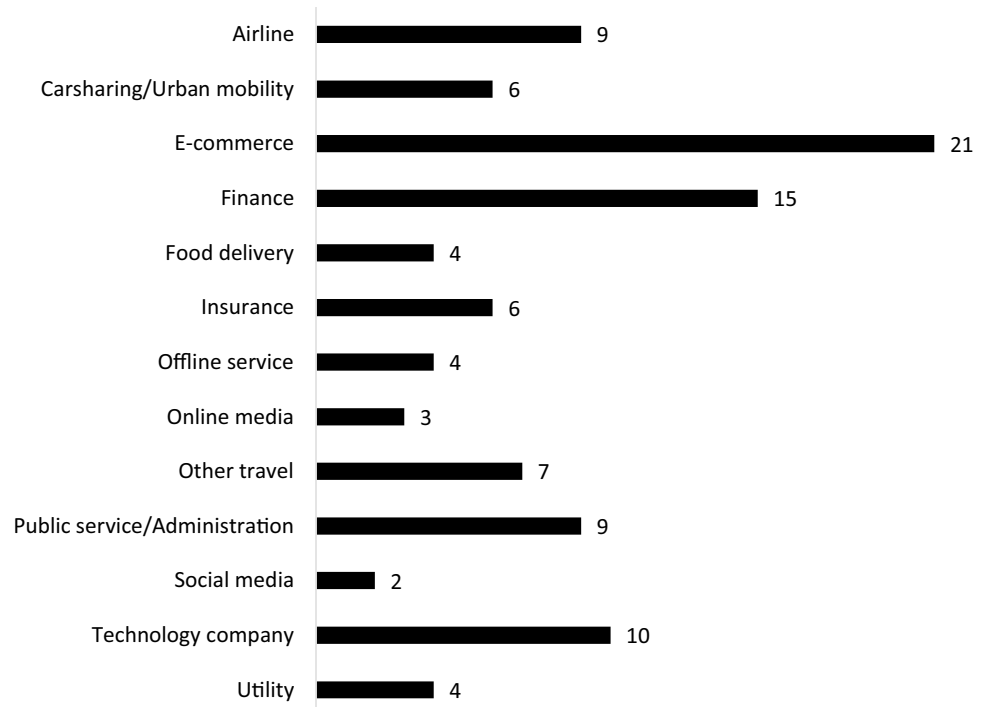


Table 2 Characteristics of the sampled companies and organizations

Legal Form	Sole Trader	German Corporation	Other German	Foreign Form	Other
	4.3%	44.7%	8.5%	42.6%	2.6%
	Min	Median	Mean	Max	S.D.
Revenue (€ million)	0.024	1,480	25,800	548,773	72,340
Employees	8	3,210	38,198	840,000	116,119
	Germany	EU	USA	Other	Global
Origin	52.7%	18.3%	22.6%	6.5%	
Target Market	32.3%	21.5%			46.2%
Publicly traded?	Yes	No			
	61.3%	38.7%			

N = 100 companies and organizations (non-representative sample)

sample. The average annual revenue was close to 26 billion Euro with a median of 1.48 billion. The average number of employees was slightly over 38,000 with a median of 3,210. German and US companies were most common, followed by EU and ‘other foreign’ (mostly Asian) businesses. UK-based companies were still counted under ‘EU’ despite Brexit. Over 60 percent of the sampled organizations were publicly traded.

In the following section, we present the results of our information requests, starting from the communications channels and the ease of contacting each organization, describing the length, complexity and transparency of the resulting replies and analyzing the substance of the reported AI applications.

3.2 Results

3.2.1 The formal aspects of requesting and receiving the AI disclosures

We begin our content analysis of the AI information request responses by describing the formal aspects of the disclosure procedures (see Table 3). In the vast majority of cases, we were able to send our requests via email or an online form, usually behind a login. In slightly over half of all cases (see ‘Ease of First Contact’), the companies had a clearly marked, dedicated communication channel for privacy requests and placing our query was easy. In another 43 percent of

the cases, a general customer support email or contact form was readily available and we could use it to pose our questions. Making the request was ‘difficult’ in a few instances where it proved challenging to find any means of contact (e.g. because the contact information was obscured by bad webpage design or the contact webpage itself was hard to locate due to poor website navigation). On average, it took a little over three weeks to get a final response, with some companies reacting within 24 hs and a few others going near or exceeding the legal deadline of three months (see ‘Reply Time’).

Research on online consumer information has shown that the sheer amount of disclosure is experienced as overwhelming and problematic by many consumers.⁵⁵ We therefore looked closely at the length of the responses. The average personalized reply to our AI information request contained 654 words (see ‘Length of Personalized Reply’). Including the respective privacy policy, which was almost always referenced in the response, the average total disclosure volume was 2,379 words (see ‘Length of AI Disclosure’). While some companies restricted their responses to 20 or 29 words to indicate no AI usage, some data-processing disclosures contained thousands of words, and one German online newspaper exceeded 19,000 words—the length of an extended scientific paper or a short novella.

Not only the length, but also the accessibility of a response has implications for consumer understanding.⁵⁶ In 60 percent of the cases, the personalized response was directly readable in the body of an email (see ‘Medium of AI Response’). A total of 24 percent of responses came by regular postal service. Only 12 percent came in some password-protected form (either as an email attachment or after logging in to the company’s online platform). Four percent of the responses led us to an open-access webpage. Thus, the majority of the responses were either immediately readable or ‘a click away’, either password-protected or freely available. The 24 responses received by the regular postal service, on the other hand, raise practical and legal questions (how difficult is it to navigate a printed response where no text search option is available; how is the right to receive an electronic response enshrined in Art. 12 GDPR safeguarded, etc.).

We next gauged the effort required to receive a response to a query about AI usage (see ‘Number of Messages Exchanged’). Sixty-four companies reacted to our very first request and provided the information we required. Another 21 required a reminder or a clarifying message before doing so. Communication with the remaining 15 companies and organizations needed a message dialogue of four to eight

messages before a final answer was provided. Such extended dialogues often occurred because companies ignored the AI focus of our inquiries and either referred us to their generic privacy policies or sent us a standard response to a GDPR subject access request. In some cases, additional messages also needed to be exchanged for the sake of authentication (see ‘Additional Authentication Requested’). Nevertheless, 88 of the data controllers in our sample did not request additional proof of identity because we included basic personal data in every request or because we had submitted our request through our password-protected account area. Additional authentication requirements included submission of a picture ID or additional customer account information and multiple authentication steps (e.g., screenshots of personal dashboard, confirmation of payment of a membership fee, etc.).

To finalize the initial description of the AI responses, we accounted for the role of the person providing the response, the degree of response personalization, and the amount and kind of personal data that accompanied the message (see the respective categories in Table 3). Most commonly, customer support personnel or data protection officers handled our requests—in 37 and 31 cases, respectively. For 24 responses the contact person’s role could not be identified. The organization’s legal counsel answered the request in eight cases. The response was ‘personalized’ in 46 cases; this means it was tailored specifically to our request and addressed the concrete points we raised substantively. Twenty-two companies sent us an ‘off-the-shelf’ GDPR response that captured the topic of AI usage, but also included much more information that we did not explicitly request. Almost a third of our sample—32 percent—took a minimal-effort approach and either copied their privacy policy in their response to us or referred us to the policy published on their website, thus providing the least amount of personalization. A total of 65 responses did not include any personal data of the account holder that they had processed; 24 responses included only the personal data categories that would or could undergo automatic processing; 11 responses included all the data categories stipulated as standard by Art. 15 of the GDPR, even though we had not requested such an extensive data disclosure.

3.2.2 The transparency and substance of the AI disclosures

We now focus on the substance of the AI disclosures we received, paying particular attention to their linguistic, visual and layout transparency and the clarity of the explanations given (see Table 4 for full details). Our yardstick for assessing disclosure transparency are the criteria published by Seizov and Wulf,⁵⁷ which are themselves based on a multidisciplinary review of best practices in disclosure

⁵⁵ Seizov et al. (n 35).

⁵⁶ Waller (n 39).

⁵⁷ Seizov and Wulf (n 42).

Table 3 Formal characteristics of the AI information request procedure

Means of First Contact	Email 81%	Contact Form 18%	Other 1%		
Ease of First Contact	Easy 54%	Neutral 43%	Difficult 3%		
Reply Time (Days)	Min 1	Median 12	Mean 23	Max 112	S.D. 28.56
Length of Personalised Reply (Words)	20	272	654.80	11,977	1,470.37
Length of Disclosure (Response + Add-ons, Words)	29	1,260	2,379.48	19,116	3,046.39
Medium of AI Response	Plain Email 60%	Secure Email 3%	Account Area 9%	Webpage 4%	Regular Mail 24%
Number of Messages Exchanged	1 64%	2 - 3 21%	4 - 5 9%	6 - 7 5%	8 1%
Additional Authentication Requested	None 88%	Picture ID 7%	Account Data 4%	Multiple 1%	
Contact Person’s Role	Customer Service 37%	Privacy Officer 31%	Legal Counsel 8%	Unknown/Other 24%	
Degree of Response Personalization	Privacy Policy Reference 32%	Generic GDPR Response 22%	Personalized Response 46%		
Personal Data Enclosed	None 65%	General GDPR Data 11%	AI-Processed Data 24%		

N=100 companies and organizations (non-representative sample)

information design.⁵⁸ The linguistic analyses, evaluations, and criticism that follow refer to the original language in which the companies provided their AI disclosures to us, i.e., either German or English. The examples we have included in the text are our own translations from the original German, unless otherwise indicated.

Only half of the responses could be deemed linguistically transparent, i.e., largely free of meandering sentences, excessive legal or professional jargon or modal phrasing that obscures the amount, kind and frequency of the automated data processing in which the organization engages. For example, the cooking kit delivery company Hello Fresh explained its automated data processing in the following straightforward manner:

We use automated decision-making (also known as profiling) to send our customers offers that are interesting and relevant to their needs. For this purpose, specific data such as the number of received boxes and the pausing of orders are processed in automatic form.

This could only possibly be improved by eliminating the passive voice in the second sentence and rephrasing it to, “For this purpose, *we use* specific data [...] and process them in automatic form.” However, the passive voice appears to be

⁵⁸ For more details on each discipline, see Seizov et al. (n 35).

used across the board, even in otherwise transparent disclosures. See, for example, the car insurance company nexible and its informative and relatively accessible reply:

As part of the price calculation, regression analyses are used. Only anonymous damage statistics are used (e.g., age or localization data). [...] No personally identifiable information is used (e.g., name, exact address, email). [...] In addition, prediction algorithms are in use for our mailbot and chatbot. These use the words you type in the contact form, chat or email as input data. For the chat, *generally* [our emphasis] no personal identification takes place. In the context of your communication with us to date, no machine learning algorithms for automated decision-making have been used.

Table 4 Transparency of the AI disclosures

	Non-transparent	Partially Transparent	Transparent	N/A
Language	18%	32%	50%	0%
Visual	7%	0%	1%	92%
Layout	8%	40%	41%	11%

N= 100 companies and organizations (exploratory sample)

The only element of linguistic doubt in the above disclosure is due to the adverb ‘generally’, which implies that personal identification may take place under specific circumstances not defined here. Aside from this lack of clarity, however, the passage uses short sentences that are easy to follow and communicate information clearly. This was not the case for 32 responses that we characterized as ‘partially transparent’, i.e., the above syntactic, linguistic, and modal problems were more pronounced. For example, the location services company Maps.me offered the following information, which we quote here in the original English:

- 3.1. In order to implement the agreement between you and us, and provide you with access to the use of the Services, we will improve, develop and implement new features to our Services, and enhance the available Services functionality. To achieve these objectives, and in compliance with applicable laws, we will collect, store, aggregate, organise, extract, compare, use, and supplement your data. We will also receive and pass this data, and our automatically processed analyses of this data to our affiliates and partners as set out in the table below and Sect. 4 of this Privacy Policy.
- 3.2. We set out in more detail the information we collect when you use our Services, why we collect and process it and the legal bases below. [A lengthy table with different data uses follows.]
- 3.3. Our legitimate interests include (1) maintaining and administering the Services; (2) providing the Services to you; (3) improving the content of the Services; (4) processing of the data that was manifestly made public by you where it is accessible by other users of the Services; (5) ensuring your account is adequately protected; and (6) compliance with any contractual, legal or regulatory obligations under any applicable law.
- 3.4. As part of maintaining and administering the Services we use the information to analyze user activity and ensure that rules and terms of use for the Services are not violated.
- 3.5. Your personal information may also be processed if it is required by a law enforcement or regulatory authority, body or agency or in the defence or exercise of legal claims. We will not delete personal information if it is relevant to an investigation or a dispute. It will continue to be stored until those issues are fully resolved and/or during the term that is required

and/or permissible under applicable/relevant law.

- 3.6. You may withdraw your consent to the collection of location data by amending your privacy settings on your device.
- 3.7. Please note, if you do not want us to process sensitive and special categories of data about you (including data relating to your health, racial or ethnic origin, political opinion, religious or philosophical beliefs, sex life, and your sexual orientation) you should take care not to post this information or share this data when using the Services. Once you have provided this data it will be accessible by other Service users and it becomes difficult for us to remove this data.
- 3.8. Please note, if you withdraw your consent to processing or you do not provide the data that we require in order to maintain and administer the Services, you may not be able to access the Services.

The sentences in this AI disclosure are mostly long and complex; they include long lists; they habitually use modal phrasing and hypotheticals (*if, and/or, may / may not*); and they spread crucial information across several sections and formats. Paragraphs 3.3., 3.5., and 3.7. exhibit most of these problems and illustrate why such disclosures are only ‘partially transparent’, language-wise. There were, however, even worse examples of insufficient clarity of information. Eighteen responses were full of overly long sentences and hazy or confusing formulations and were thus classified as ‘non-transparent’. Consider the response from Booking.com, which we quote here in the original English:

When you make calls to our customer service team, Booking.com uses an automated telephone number detection system to relate your telephone number to your existing reservations—this can help save time for both you and our customer support staff. Not all calls are recorded and recordings are kept for a limited amount of time and automatically deleted thereafter, unless Booking.com has a legitimate interest to keep such recording for a longer period, including for fraud investigation and legal purposes.

Booking.com accesses communications and may use automated systems to review, scan, and analyse communications for security purposes; fraud prevention; compliance with legal and regulatory requirements; investigations of potential misconduct; product development and improvement; research; customer engagement, including to provide you with information and offers that we believe may be of interest to you; and

customer or technical support. We reserve the right to block the delivery of or review communications that we, in our sole discretion, believe may contain malicious content, spam, or may pose a risk to you, accommodation partners, Booking.com, or others.

If we use automated means to process personal data which produces legal effects or significantly affects you, we will implement suitable measures to safeguard your rights and freedoms, including the right to obtain human intervention.

The text above is non-transparent in a number of ways. The sentences are mostly long. Many of them tend to make one unclear statement and proceed to qualify it with conditions, which result in even less clarity. For example:

‘Not all calls are recorded [*What are the selection criteria?*] and recordings are kept for a limited amount of time [*What is the amount of time?*], unless Booking.com has a legitimate interest [*What constitutes a legitimate interest?*] to keep such recording for a longer period [*How long is that? What is the maximum period?*], including for fraud investigation and legal purposes [*What other ‘legal purposes’ could be legitimate here?*].’

Thus, the first paragraph only appears to answer data privacy questions, while in fact it raises additional ones. The long middle paragraph presents a cavalcade of *possible* uses of automated data processing that mix legal (e.g., compliance), criminal (e.g., fraud prevention), and commercial (e.g., product development) purposes indiscriminately. The final paragraph promises the implementation of ‘suitable measures’, which remain unspecified. In combination, these

and similar linguistic choices make disclosures such as this one non-transparent.

In terms of visual transparency, the majority of the AI responses (92 percent) did not use any visual elements. Of the remainder, seven used rudimentary visual signage (e.g., an additional font colour, underlining, arrows, etc.) that was not consistent or extensive enough to substantially affect transparency. One company, on the other hand, relied on more advanced visualization by introducing different topics with thematic icons and by highlighting important passages consistently. In general, the call by Berger-Walliser and colleagues for greater integration of verbal and visual language in legal documents remains unheeded.⁵⁹

Turning to layout, 41 percent of the companies achieved transparency by breaking down their AI disclosures into thematic paragraphs and giving each paragraph a meaningful heading that made the hierarchy and structure of the text as a whole apparent. Another 40 percent of the responses were ‘partially transparent’, meaning that the text was broken down into more or less palatable chunks, but the themes were not as clearly identifiable and helpful headings were missing. Eight responses got a ‘non-transparent layout’ rating because they did not have a discernible text hierarchy and structure. Eleven responses consisted of a single paragraph and thus had no layout qualities to rate; typically, such brief and amorphous responses either failed to report any automated decision-making or directed us straight to the official privacy policy.

Overall, 34 percent of the companies and organizations we sampled did not disclose any AI usage. We then classified the different applications of automated data processing reported by the remaining 66 companies into five broad categories (see Fig. 3).

Fig. 3 Types of automated data processing reported by $N=66$ companies and organizations that reported at least one form of AI-supported data processing (exploratory sample). The sum of percentages exceeds 100 because many organizations reported more than one item



⁵⁹ Berger-Walliser et al. (n 41).

Table 5 Assessment of the AI disclosures

Relative to GDPR	Non-Transparent and Insufficient 5%	Transparent but Insufficient 21%	Non-transparent but Sufficient 17%	Transparent and Sufficient 57%
Relative to respondent's own privacy policy	Response Contradicts Policy 1%	Response Is Less Detailed than Policy 35%	Response Repeats Policy 31%	Response Is More Detailed than Policy 33%
Overall, AI usage is...	Not Present 31%	Opaque 6%	Not Fully Addressed 45%	Fully Addressed 18%

N = 100 companies and organizations (non-representative sample)

Artificial intelligence was most frequently used (68 percent of positive AI responses) for the purpose of improving products and services, e.g. by tracking customer behaviour and product or service usage. AI-assisted credit scoring was mentioned in one-third of the positive AI responses, mostly in the financial, insurance and e-commerce sectors where creditworthiness plays a major role. Profiling, i.e., “automated processing of personal data [with the purpose of] evaluating certain aspects relating to a natural person” and making predictions about them (Art. 4, para. 4 GDPR), was mentioned in 30 percent of the positive AI responses, and automated decision-making in 26 percent. Cookies (mentioned in 55 percent of the positive AI responses) are a grey zone: although every website uses them nowadays, some organizations did not count them as part of automatic processing while others did, hence the surprising outcome that two-thirds of our sample did not report on using this popular automatic tracking mechanism.

3.2.3 An overall assessment of the AI disclosures

Finally, we produced an overall assessment of the AI disclosures we received (see Table 5). Based on the specific criteria we have covered in this section, we determined to what extent the responses fulfilled the requirements of the GDPR, how they compared to the official privacy policy of each company and what the extent of AI usage reported actually was. According to the standard for the provision of information laid down in the GDPR, set out in Art. 12(1) GDPR as being of a “concise, transparent, intelligible and easily accessible form [and] using clear and plain language”, 57 percent of the responses were sufficiently clear and comprehensive. Twenty-one percent were transparent but did not comply (e.g., because they failed to provide crucial details on the nature or purposes of automated processing). Seventeen percent were not transparent (e.g., due to linguistic or layout shortcomings) but provided all the necessary information. The remaining five percent missed the mark in regard to both transparency and informativeness. Therefore, just over half of the

responses passed the ‘GDPR test’, and yet a substantial minority lacked the necessary information or clarity, or both.

We also compared the AI disclosures with the information contained in the privacy policies. The three most common forms of deviation we identified (response was more or less detailed than the published policy or just as detailed as it) occurred at roughly identical frequencies. One company, an online shop, provided a response that contradicted its privacy policy. The response claimed that third-party cookies were the only example of automated data processing, while the privacy policy asserted that the online shop uses automated data processing in two instances. The first of these is not specified and the second is automated message filtering for spam and scam detection.

As regards the substance of the AI disclosure, 31 companies stated that they do not employ automated decision-making, and six did not provide a clear answer, possibly concealing the amount and nature of such processes, with the result that their AI disclosures are rather opaque. A typical example of opacity came from the online gaming platform Steam (quoted here in the original English):

Steam does not perform any data processing unless you're actively using a feature that requires that processing.

Additional information can be found in the Privacy Policy Agreement.

In addition to the circumlocutionary nature of the personalized response, the linked Privacy Policy Agreement did not explicitly mention automated decision-making or data processing. When we contacted the customer support team with a request for clarification this was not helpful:

The previously linked privacy policy contains Steam's official statements regarding this matter.

As an example, Steam Store recommendations are only processed when accessing the Steam Store. If you wish to avoid having that data processed, you can set Steam

Table 6 AI disclosures of the airlines and frequent-flyer programmes in our sample

	Cookies	AI for service optimization	AI for consumer profiling	AI for automated decisions	AI for credit rating
Air Baltics	✓	✗	✗	✗	✗
British Airways	✗	✗	✗	✗	✗
Delta	✗	✗	✗	✗	✗
EasyJet	✓	✗	✗	✗	✗
Finnair	✗	✗	✓	✗	✗
Lufthansa	✗	✗	✗	✓	✓
Miles & more	✗	✓	✗	✗	✗
Ryanair	✗	✗	✗	✓	✗
Turkish Airlines	✓	✓	✓	✗	✗

to open directly to your library rather than the Steam Store.

While not as opaque as the above example, an additional 45 responses also failed to provide a full explanation of their AI data processing (for example, they mentioned automated services or processes but did not reveal the mathematical or other logic behind them). Such responses stated that the company uses automated processing for one or more purposes but included no further details about the process or its practical and legal implications for the consumer. Only 18 companies offered a satisfactory disclosure that covered the amount, kind, mathematical nature and consumer implications of their automated data processing activities.

3.2.4 Case studies

Our findings, particularly the ones relating to the transparency of the AI disclosures (see Table 4) and the different uses of automated data processing (see Fig. 3), prompted us to take a closer look. Given the widespread adoption of AI in a number of industries and the great potential that such technologies hold, the overall low numbers of automated processing the companies in our sample reported made us look into several cases by way of example.

3.3 Airlines and frequent-flyer programs

We had nine airlines in our sample. Although their business models and their processes are likely to be nearly identical, their AI disclosures included vastly different information. We present a brief yet telling comparison in Table 6.

British Airways and Delta did not report any automated data processing, while Air Baltics and EasyJet only admitted to using cookies. Finnair reported profiling consumers and placing them in a ‘customer segment’ based on their travel history. Ryanair’s disclosure included the description of an automated seat assignment algorithm designed to keep the

machine in balance. For its frequent-flyer programme Miles & More, Lufthansa stated that automated data processing is employed to improve advertisement targeting. Lufthansa’s own AI disclosure, on the other hand, mentioned creditworthiness checks and an automated process for detecting fake or stolen credit cards or other payment methods. Finally, Turkish Airlines (after several email exchanges back and forth, initially in Turkish) listed cookies, service optimization and profiling as uses of AI, without any additional information.

The heterogeneity of the disclosures of the eight airlines and one frequent-flyer programme is indicative of a lack of a common standard on AI reporting in an industry that is highly datafied,⁶⁰ offers highly standardized products and relies on streamlined business processes. In fact, the whole truth would likely be achieved by combining *all* nine disclosures. The websites of all nine companies used cookies as of May 2020. They gather and analyze personal data automatically in order to improve the customer experience and to hone their marketing efforts. Along the way, they profile customers and quite possibly engage in price discrimination.⁶¹ A seating algorithm that maintains the airplane in balance hardly sounds like a single airline’s proprietary invention and is almost certainly in use across the industry. The same would also be true of preventing credit card fraud and making sure customers will pay their fares.

3.4 Insurance

Although the insurance industry has been late to adopt digital trends compared to others (Deloitte 2017), it is catching up quickly. In the last few years, automated processes have been implemented both in customer support and in claims management, risk assessment and best-fit insurance policy

⁶⁰ For a general discussion of datafication, see Van Dijck (n 3).

⁶¹ Van Boom et al. (n 25).

Table 7 AI disclosures of the insurance companies in our sample

	Cookies	AI for service optimization	AI for consumer profiling	AI for auto-mated decisions	AI for credit rating
Arag SE	✗	✓	✗	✓	✓
Europa-GO	✓	✗	✗	✗	✗
Nexible	✗	✓	✗	✓	✗
R + V	✗	✗	✗	✗	✗
Techniker Krankenkasse	✗	✗	✗	✗	✗
HUK-Coburg	✗	✗	✗	✓	✓

comparisons.⁶² The OECD highlights the increasing implementation of risk-based insurance pricing since “more data could improve the predictability of policyholder behaviour or incidents”, with a particular focus on “risk sensitivity or propensity to switch”. The result is a finely tuned, data-driven, and more individual risk assessment.⁶³

If we turn to the AI disclosures that we gathered in the present study, many of the insurance companies in our sample do not yet seem to have embraced digitalization or advanced AI (see Table 7). In fact, neither leading German general insurer R + V nor health insurer Techniker Krankenkasse give information on any automation whatsoever, and car insurer Europa-GO listed cookies as the sole source of automatic data processing. The other car insurance company in the sample, nexible, stated that it uses AI for chatbots and mailbots (i.e. for customer service optimization) and for pricing (i.e. regression analyses based on risk factors such as age, primary address, insurance history and other details which were not named). A subsidiary of the insurer HUK-Coburg mentioned automated pricing, risk determination and decisions on creditworthiness, as did Arag SE, which also uses AI to improve its marketing and customer-care offerings.

As with airlines, it is most likely in reality that AI usage is much more widespread in the insurance industry than the disclosures revealed. Again, the problem is most likely not intentional obfuscation but rather a lack of clear standards for communicating information. Given the expectations of “an explosion of data from connected devices” and ever-growing avenues for risk minimization and customer interaction,⁶⁴ such standards will need to be formulated and enforced much more clearly in order to capture the rapidly increasing complexity and reach of automated data processing in this branch of industry.

3.5 Carsharing/Urban Mobility

The one-way carsharing and urban mobility industry has experienced a boom thanks to the success of several industry leaders, for instance Daimler’s car2go and BMW’s DriveNow, which merged in 2019 to form the common platform ShareNow,⁶⁵ as well as rising stars, such as Volkswagen’s all-electric venture WeShare. Tracking each vehicle’s route, predicting common paths, and maintaining a uniform distribution of vehicles across the service area are crucial elements of a carsharing business model.⁶⁶ Vetting each customer’s behavioural and financial good standing and keeping track are equally important tasks. Ride-hailing and ride-sharing services like Uber and FreeNow have similar data practices, as does the parking app Easypark. It is therefore reasonable to expect that a significant amount of automated data processing is carried out in this business sector. Their disclosures (see Table 8) are, however, rather meagre.

Uber is the only urban mobility company that met our expectations in regard to details on the amount and kind of automated processing. Its disclosure included cookies that track usage, automated service optimization (e.g. removal of drivers and delivery partners with poor ratings, matching customers and drivers based on ‘availability, proximity, and other factors’), automated user-profiling to customize incentives and automated decision-making for dynamic pricing for rides and deliveries (i.e. a price-formation algorithm) as well as banning users whose behaviour is automatically classified as fraudulent, dangerous or harmful. BlaBlaCar also mentioned cookie use and automatic service optimization (i.e. improved targeting for marketing purposes). FreeNow uses automated decision-making to regulate the modes of payment a customer is offered. The company provided a well-formulated explanation of how the algorithm (of the ‘random forest’ type) determines whether a customer is allowed to make payments via the app or whether only cash or debit card payment is possible. Easypark merely

⁶² Thomas (2020).

⁶³ OECD (2020).

⁶⁴ Balasubramanian et al. (2018).

⁶⁵ Maslen (2019).

⁶⁶ Arakawa (2017), Enzi et al. (2020).

Table 8 AI disclosures of the carsharing and urban mobility companies in our sample

	Cookies	AI for service optimisation	AI for consumer profiling	AI for automated decisions	AI for credit rating
BlaBlaCar	✓	✓	✗	✗	✗
DriveNow	✗	✗	✗	✗	✗
Easypark	✓	✗	✗	✗	✗
FreeNow	✗	✗	✗	✓	✗
WeShare	✗	✗	✗	✗	✗
Uber	✓	✓	✓	✓	✗

declared its use of cookies and no automated data processing of any kind. DriveNow (before merging with car2go to form ShareNow) and WeShare declared that they used no AI algorithms, cookies or other kind of automated data processing whatsoever—a rather unexpected outcome that raises further questions about the disclosure standards the GDPR is meant to establish and maintain.

3.6 Credit scoring and creditworthiness

The use of computer-assisted creditworthiness checks (or *Bonitätsprüfung* in German) has been growing since the 1980s.⁶⁷ Nowadays, advanced, fully autonomous AI systems evaluate consumers' creditworthiness and play an important role in online trade, particularly in the finance, insurance and utilities sectors and also in many e-commerce transactions, i.e. in the sale of both goods and services. Neural networks seem particularly popular in this area,⁶⁸ and the growing wealth of electronic data available to companies has led researchers and practitioners to include all kinds of additional information in the creditworthiness calculations, even consumers' social media information and interactions.⁶⁹ Dorffmeister points out that EU law encourages thorough creditworthiness checks but provides for little or no oversight or disclosure requirements in regard to their design and implementation.⁷⁰

A total of 22 companies in our sample drawn across several industries disclosed that they use automated creditworthiness checks. Most did not provide any details of the process employed, which can be expected in the German context, at least to a degree. In a landmark case in 2014, the German Federal Court ruled that the creditworthiness algorithms of the leading German credit scoring agency SCHUFA could remain hidden, confirming the right of firms

to keep trade secrets.⁷¹ In 2018, a consumer demanded that a company send him his full credit scoring report, including the partial scores that made up his total creditworthiness rating, but the Regional Court of Wiesbaden did not find his demands reasonable.⁷² In the German jurisdiction, companies that engage in automated creditworthiness checks are therefore free to reveal little or none of the processes involved without fear of legal repercussions. Our consumer survey showed that the majority of the respondents would nonetheless like to have detailed information about automated decision-making, including the personal data employed and the basic logic that drive algorithmic analyses and decisions. We present the 22 AI disclosures we received in Fig. 4.

The typical response on creditworthiness checks included a basic notice that the automated check was in place and that further details could not be provided. SCHUFA, the most popular credit scoring company, merely disclosed that it uses “modern mathematical statistical procedures” to evaluate each person. Input variables, weightings, and all other details were deemed to be trade secrets. Many companies who use the services of SCHUFA, Infoscore Consumer Data GmbH or similar services used a similar disclosure strategy. For example, Lufthansa declared that it uses an automated validity, liquidity and legality check on its customers' chosen payment methods to prevent bank fraud and stated that “understandably” it cannot offer information on the logic of the algorithm employed for security reasons. Advanzia Bank stated that it employs credit scoring algorithms to automatically determine credit card limits and also asked customers to understand that it cannot share the logic involved. Another major German online bank assured us that the creditworthiness check “rests on a mathematically and statistically recognized procedure that is tried and tested”, but did not

⁶⁷ Kulmann and Reucher (2000).

⁶⁸ Dittmar and Hilbert (2015).

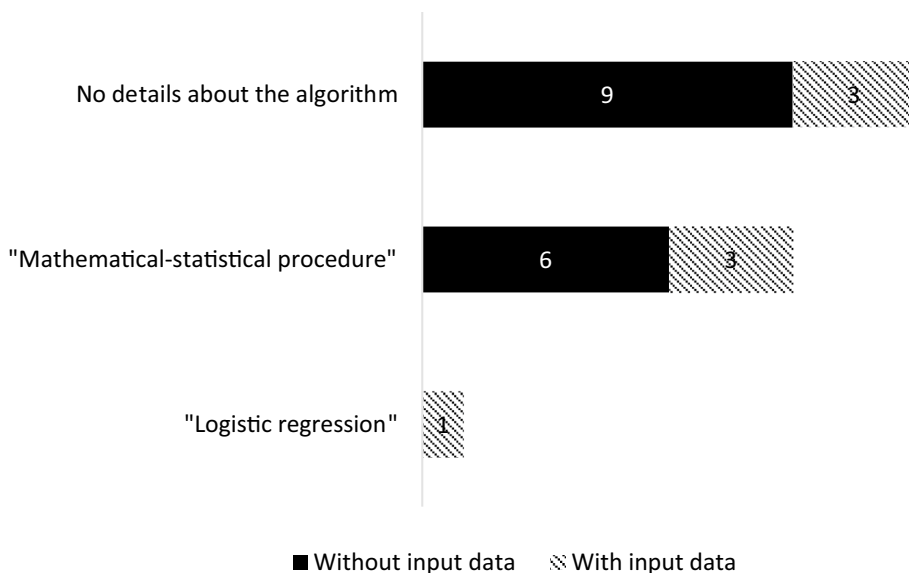
⁶⁹ Mengelkamp (2017).

⁷⁰ Dorffmeister (2017).

⁷¹ *Umfang des Auskunftsanspruchs gegen die Schufa-Scorewerte*, Case VI ZR 156/13, 747 NVwZ (BGH [German Federal Court] 2014).

⁷² *Case 5 O 214/18*, 33,343 BeckRS (LG Wiesbaden [Regional Court of Wiesbaden] 2018).

Fig. 4 Types of AI-assisted creditworthiness assessments. $N=22$ companies disclosed their use of AI-assisted credit scoring (exploratory sample). Only one company disclosed both the method and the categories of the personal input data used in the calculation. Nine companies referred to a “mathematical-statistical procedure” and three of them also specified the categories of input data employed. The remaining 12 companies did not provide any details on their credit scoring algorithms, and nine of them also failed to disclose details on the input data



offer any more details. The online shop of Saturn responded in similar terms.

At the other end of the spectrum there was a minority of companies who offered more detailed explanations of their automated creditworthiness checks, which invariably included the use of regressions. In the case of the online shop MisterSpex, AI-supported credit scoring is used only for consumers who wish to pay per invoice or wire transfer. According to the disclosure, the check consists of a logistical regression computation that takes address and past payment behaviour as its independent variables to estimate the likelihood of payment issues. A comparison with the behaviour of consumers with similar characteristics also influences the final estimate. Vorwerk, the manufacturer of the popular kitchen robot Thermomix, also made a more detailed disclosure. It included the personal data categories that go into the “mathematical statistical procedure”, the result of which is not “a mere number but a detailed recommendation about the appropriate payment methods” to which the consumer should be granted access.

Overall, MisterSpex was the only company that came close to meeting the majority of our survey respondents’ expectations of transparent AI disclosure. Although it did not go into the underlying logic of the process, it revealed both the kind of statistical procedure and the personal data fed into the automated creditworthiness check it conducts. The remaining 21 businesses, many of them banks and insurance companies, revealed little or nothing about the procedures and input variables used in their creditworthiness evaluations and thus failed to meet the expectations of nearly four-fifths of consumers, according to our survey.

4 Conclusions

In this section, we revisit our guiding questions to draw the relevant conclusions and shape the policy debate.

4.1 How do consumers perceive AI usage and how do they wish to be informed about it?

Our consumer survey revealed that many consumers are not yet well versed in AI technology and have not made many conscious experiences with it. In addition, 90 percent of them had never made use of their GDPR-based right to information and explanation, either because they did not know about it or because they did not feel the need to do so. On the other hand, a majority of the respondents insisted on human review of automated decisions, even if they had not experienced any negative consequences, and also demanded extensive AI disclosures, including statements on algorithmic logic, personal data, and sometimes even the source code. These are requirements that go beyond the GDPR’s provisions. Overall, there is a gap between consumers’ wishes for information and their knowledge of how to obtain it. Almost 80 percent of the respondents wanted a substantive disclosure on AI, and the majority of them expected information not only about the logic but also on the personal data that go into the decision-making process. This finding echoes the recommendation of the IEEE that the basis of every algorithmic decision be clearly identifiable.⁷³ The results of our content analysis show that a mere 18 per cent of the responses met that standard and 51 per cent were either opaque or did not address AI usage fully

⁷³ Shahriari and Shahriari (n 5).

(see Table 5). Furthermore, 65 per cent of the responses did not include any mention of personal data, thus leaving any AI explanations vaguer than most consumers would expect.

4.2 How is AI currently used by companies operating in Germany?

Based on the self-reports, 66 percent of our sample of companies use some form of AI-supported automated data processing (see Fig. 3). Most of them use AI to improve the customer experience and to optimize product and service offerings. About half of those companies say that they use cookies and one-third use credit scoring. At 30 and 26 percent, respectively, profiling and automated decision-making are the least common usages of AI. However, there is reason to believe that profiling, automated decision-making, and credit-scoring are not declared as often as they are used. Both academic and industrial research point in the direction of much more widespread and advanced algorithmic data processing, especially in technology- or digitization-heavy industries such as banking, insurance, air travel, and e-commerce. Visits to each of the companies' websites confirmed that all of them use cookies, but only 55 percent included cookies in their AI disclosures. If businesses are unsure about whether or how the usage of a relatively simple automated processing mechanism such as cookies should be declared, there is reason to doubt the accuracy and completeness of the information provided on other, more advanced AI applications. According to our case studies, companies and organizations habitually declared lower use of AI than could be reasonably expected. Thus, according to the self-disclosures, AI is not yet widely used by the 100 companies operating in Germany that we sampled, but this conclusion should be taken with a big grain of salt.

4.3 Is the GDPR's right to AI disclosure effective in informing consumers?

As previously pointed out by legal commentaries,⁷⁴ the findings from our content analysis of the disclosures on AI use confirm that the requirements of the GDPR are too vague and leave substantial room for interpretation. The result was a collection of highly heterogeneous disclosures which, as we point out above, are likely incomplete. Even if we leave our concerns regarding completeness aside, there are wide variations in disclosure practices. These relate to how easy it is to make a formal request for information, the number of messages we had to send before we received a meaningful response, the authentication requirements, the length of the disclosure, and the amount and nature of the personal data

included in each AI disclosure (see Table 3). While scholars in the field of information obligations have repeatedly called for both standardization⁷⁵ and simplification,⁷⁶ the GDPR appears to fail to deliver on either of these requirements. Turning to the substance of the AI disclosures, they either described automated data processing in an 'opaque' fashion or did not address the matter fully in 51 percent of cases; a mere 18 percent of disclosures were sufficiently thorough (see Table 5). The case studies that we presented in Sect. 3.2.4 further underscore the variations both between companies within the same industry and between branches of industry. Taken together, these findings raise concerns about the effectiveness of the GDPR-mandated AI disclosures, both in terms of the data subjects' ease of access to information and of the quality and completeness of the information provided.

4.4 Are the GDPR disclosures transparent?

Finally, the transparency of information notices has been a highly controversial topic that scholars have approached from a variety of angles.⁷⁷ We evaluated the transparency of the AI disclosures according to two standards: (a) the requirements of the GDPR and (b) the recommendations made by Seizov and Wulf.⁷⁸ A little over half of the disclosures met the GDPR's relatively lax and general standards and we deemed them 'transparent and sufficient' (see Table 5). Not being able to meet even such a generic and forgiving standard of transparency does not inspire too much confidence. When we considered the language, visual design and layout of the AI disclosures as recommended by Seizov and Wulf (see Table 4), we came to similar conclusions. The language was transparent in merely half of AI disclosures, visuals were seldom used and were transparently employed in only a single instance, and the layout was more likely to be non-transparent or partially transparent (48 percent total) than fully transparent (41 percent). Therefore, a significant fraction of the AI disclosures failed to meet either the modest transparency requirements of the GDPR or Seizov and Wulf's specific transparency recommendations. The need for improvement in this area is apparent.

5 Discussion

Our findings identify a mismatch between consumer needs and expectations and the legal and practical status quo with regard to the provision of information to consumers on the

⁷⁴ Feiler et al. (n 16); Gierschmann et al. (n 30).

⁷⁵ Ben-Shahar and Schneider (2014).

⁷⁶ Elshout et al. (n 20); Marotta-Wurgler (n 21).

⁷⁷ See Seizov et al. (n 35) for an interdisciplinary overview.

⁷⁸ Seizov and Wulf (n 42).

AI-assisted automated processing of personal data. The likely reason for this mismatch is the topicality of the subject of AI regulation and the long drafting process of the GDPR, as a result of which policy makers were unable to take into account the importance that AI regulation would have for consumers in the future. The trailblazing nature of the GDPR and the fierce lobbying for and against this new regulation by opposing interest groups led to a drafting process that spanned a whole decade, beginning with the initial policy work and ending when the law came into effect. During this period, public and policy discourses were dominated mainly by data privacy and surveillance concerns, due not least to the Snowden revelations.⁷⁹ While the topics of AI data processing and decision-making were already addressed in these discourses, they had not yet reached the significance that they have today. Recently, public awareness for AI regulation has increased because of the steady growth of AI-powered consumer products and services and prominent scandals such as the Facebook–Cambridge Analytica scandal and critical media coverage on, for example, the first (fatal) crashes involving driverless cars or the development of AI powered military equipment. Thus, it is only recently that the European Commission⁸⁰ and other international policy makers⁸¹ have reacted and put the topics of AI ethics, control, and oversight higher on their political agendas. This increase in attention came too late to have an effect on the rules applying to AI in the GDPR, which had by then already been adopted.

Therefore, although the GDPR is the most developed piece of legislation on data protection currently in force, the guidance that it provides on how companies should inform data subjects about AI and automated decision-making appears to be too superficial and open to interpretation. Our empirical results suggest that the Regulation fails to lead to the production of reliable and comparable disclosures that would satisfy consumers' expectations. And this is despite the publication of two extensive additional pieces of guidance in 2017, the year after the adoption of the GDPR,⁸² by the Article 29 Data Protection Working Party of the European Data Protection Board. The first of these, referring specifically to Article 22 of the GDPR, are the 'Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679' (WP 251). The

Guidelines aim to clarify several AI-related provisions of the GDPR, defining key concepts and offering further interpretations of the regulation. Yet the document appears "more to be making or extending rules than to be interpreting them" and provides "only partial clarity—and perhaps even some extra confusion."⁸³ Secondly, the 'Guidelines on Transparency under Regulation 2016/679' (WP 260) provide extensive interpretative background especially to the requirement by Article 12 GDPR that all communication with the data subject of the type that we asked for in our information requests be of "a concise, transparent, intelligible and easily accessible form, using clear and plain language". As we have shown, many data controllers fail to comply with this rule, which implies that any additional clarification provided by the Guidelines has done little to reduce the gap between consumer expectations and corporate disclosure practices.

The GDPR's lack of punch is also evidenced by the widespread under-reporting of the use of AI by the businesses we contacted, for example in the airline, insurance, urban mobility, and credit scoring industries. There seem to be multiple reasons for the problem of under-reporting. Many businesses appear to deal with inquiries about AI usage only as part of general data subject access requests. While we specifically inquired about AI usage, almost a third of all companies we contacted needed additional prompting to be able to process our requests for clarification. Furthermore, many of the responses we received contained only snippets of information on AI use copied from the businesses' general data privacy policies or their standard, preformulated data access request responses. From this practice we infer that, while data subject access requests are common and most companies have established internal business processes to respond to them swiftly and relatively comprehensively, enquiries about AI usage are currently not specifically addressed, or if they are, only as an afterthought to the former. From an enforcement point of view, apart from those cases that we have cited, we are not aware of any recent major German court cases or notable measures taken by German data protection authorities involving Art. 15, para. 1(h) or Art. 22 GDPR.⁸⁴ Taken together, the modest and rather vague regulatory framework of the GDPR, the fact that consumer interest in automated data processing has only recently begun to emerge, and the lack of enforcement of the relevant articles in the GDPR seem to create an environment in which businesses do not have sufficient incentive to provide comprehensive information on their use of AI. While it is beyond the scope of this paper to discuss solutions for all factors contributing to this problem, we now review how standardising GDPR-mandated AI disclosures could improve the way

⁷⁹ Laurer and Seidl (2021).

⁸⁰ High-Level Expert Group on AI (2019).

⁸¹ See, for example, UNESCO (2019), OECD AI Policy Observatory (2019).

⁸² A remarkably short response time. In the context of European consumer contract law, the European Commission only recently started to elaborate, in a "Commission Notice", a more specific explication of the transparency requirements contained in a Council Directive adopted 26 years earlier. See European Commission (2019).

⁸³ Veale and Edwards (2018).

⁸⁴ See Wulf and Seizov (n 7) for a review of the applicable case law.

companies inform data subjects about AI and automated decision-making. Standardisation is a promising avenue to improve information disclosure in general,⁸⁵ and we think it would also work well in regard to AI algorithms, which require consumer information disclosures of a type that are among the most complex.

Wulf and Seizov propose one standard solution for the disclosure of AI algorithmic data processing.⁸⁶ According to their proposal, companies would have to disclose the purposes and basic logic of the automated processing and illustrate how the algorithm classifies consumers by providing representative examples of consumer clusters and their characteristics. In addition to the data that feed into the algorithm, the business should supplement those user categorizations with additional business intelligence insights⁸⁷ that describe the different groups more thoroughly and make the reasoning behind the classifications even clearer. This explanation of AI processes should take the form of a standardised one-pager and be regularly updated every time the business changes its automated processing practices. Combining technical and business intelligence explanations helps make this AI one-pager both detailed and sufficiently accessible. The transparency of AI disclosures can also be enhanced by adding a more sophisticated form of legal document design that moves beyond the sole use of the written word. Berger-Walliser et al. make a convincing case for introducing visuals into legal texts.⁸⁸ As much of modern culture and discourse are turning to visual and multimodal communication,⁸⁹ law and legal documentation should follow suit. By combining design thinking and traditional legal document principles, Berger-Walliser et al. develop a five-step model for enhancing legal documents with pertinent forms of visualisation that achieve effective, empathetic and targeted communication with a multitude of audiences.⁹⁰ A combination of an illustrative, legal document design and a standard AI one-pager would then stand a good chance of realising the high goals that the GDPR sets out in its recitals we quoted above, but then fails to achieve due to the low level of specificity of Articles 15 and 22.

According to the results of our study, this approach would satisfy the information requirements of 85 percent of the consumers we surveyed. However, 15 percent of the respondents asked for even more detailed information than this proposal provides for. One solution that might also satisfy these consumers' information needs would be to

require AI operators to provide two separate disclosures. The standardised one-pager that Wulf and Seizov are proposing could cater to the information requirements of the typical consumer. A separate, more advanced AI disclosure with additional technical information such as the source code of an AI application would be directed only at technologically educated readers who have sufficient time and resources to engage with the AI-powered product or service in an in-depth manner. This audience would include sophisticated consumers and other stakeholders such as consumer organizations, privacy advocates and activists, supervisory authorities and policy makers. Such a dual disclosure strategy might then also indirectly benefit the typical consumer who does not read the advanced AI disclosure him or herself. Some professional readers may act as information intermediaries between the AI operator and the typical consumer. They could inspect in detail, for example, the computer code and algorithms that power the AI product or service and then communicate their findings in a plain manner to a lay audience. Yet other professional readers might monitor the compliance of AI operators with the GDPR and future regulation in this field. Such activities could promote enforcement and thereby help overcome the problem of the widespread under-reporting of the use of AI that we have identified. However, this dual disclosure strategy would not be without a risk. It would require that the advanced AI disclosure be easily findable yet clearly separated from the AI one-pager, so that average consumers bypass it while professional readers can still locate it without trouble. Otherwise, this approach would do more harm than good insofar as it would overload the typical consumer with too much information that they would be unable to process.⁹¹

Finally, the question remains as to how these calls for increased AI transparency can be implemented. Above we have pointed out that the regulatory framework of the GDPR is only modestly effective when it comes to automated data processing and also that the European Commission is becoming increasingly interested in ensuring that regulation of AI applications is ethically driven. In the long-term, we may hope for tighter AI regulation to complement the incomplete and vague rules of the GDPR in this domain, which the two guidelines we discussed do not appear to have substantiated much.

In April 2021, the European Commission has already presented a draft Artificial Intelligence Act.⁹² Compared

⁸⁵ See, for example, Luzak (n 35).

⁸⁶ Wulf and Seizov (n 7).

⁸⁷ See, for example, Camilleri (2018).

⁸⁸ Berger-Walliser et al. (n 41).

⁸⁹ Bateman et al.; Seizov and Wildfeuer (n 40).

⁹⁰ Berger-Walliser et al. (n 41).

⁹¹ Ben-Shahar et al. (n 74), pp 36–37 and 185–190.

⁹² 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.

to the GDPR, the draft AI Regulation contains additional transparency and information requirements, particularly for high-risk AI systems. However, the draft AI Regulation again falls short of providing a concrete yardstick of *how* operators should fulfil these transparency and information requirements. Instead, it once more relies on generic formulations and vague terms such as the stipulation that high-risk AI systems shall be “sufficiently transparent” and that information disclosures explaining the various characteristics of these systems shall be “concise”, “complete”, “clear” “relevant” and “comprehensible”, without defining what these terms actually mean in this context. On a theoretical level (“law in the books”), the draft AI Regulation thus provides an improvement over the GDPR’s modest regulatory framework for AI systems. Given the vagueness of the draft Regulation and the room it leaves for interpretation, our empirical results suggest that it may do little to improve the regulatory standard (“law in action”). Thus, the expectations and needs of the data subjects may still not be met. This concern is shared by the European Data Protection Board and the European Data Protection Supervisor, who state in their joint opinion on the draft Regulation that “[e]nsuring transparency in AI systems is a very challenging goal (...). The Regulation should promote new, more proactive and timely ways to inform users of AI systems on the [decision-making of these systems]”.⁹³ In the context of the draft AI Regulation, transparency-enhancing measures along the lines that we have proposed above are thus particularly relevant. We hope that the European Commission will consider these policy suggestions when finalizing its legislative proposal.

Until such time as such transparency-enhancing measures are enacted, there are various ways in which disclosure practices could be improved under the current framework of the GDPR. All of them are, however, far from optimal. Short of passing a new regulation, one option is to rely on the courts to further develop the law in this field. However, in the past, courts in the German jurisdiction have not been particularly active in promoting consumer rights in the area of automated data processing.⁹⁴

Above we have also reviewed recent case law from the Dutch jurisdiction that paints a similarly mixed picture of the degrees of restrictiveness of Dutch courts when it comes to granting data subjects access to information about the algorithms that affect them. The Amsterdam District Court granted data subjects access to this information in only one out of three of the cases in the review. Here the data subjects had to convince the court that their rights were significantly negatively affected by a fully automated algorithm

that processed their data without human intervention.⁹⁵ The data subjects in this case were professionals (taxi drivers) and they were affiliated with (and possibly also supported by) a trade union. It is unlikely that data subjects not acting in their professional capacity but as consumers would regularly be able and willing to provide such evidence. Furthermore, in one of the judgements by the Amsterdam District Court it appears from the text of the decision⁹⁶ that it was not until this was established during the court proceedings that the data subjects discovered that the decision-making about which they were complaining was not fully automated. Because the data controller (i.e. Uber) was able to demonstrate that, contrary to its own privacy policy, human oversight was in fact involved in the significant negative decision, the request for access to information about the algorithm employed was denied. This demonstrates, in line with the results of our study, how inaccurate and non-transparent the information that companies provide to consumers about their use of AI algorithms can be.

It is probably too optimistic to expect new and comprehensive case law to emerge in the near future that would fundamentally change the current status quo. On the contrary, in the past German courts have not exactly been famed for their contribution to fostering transparent consumer disclosures. We cited above a landmark case from 2014 in which the German Federal Court ruled that the creditworthiness algorithms of the leading German credit scoring agency could remain hidden from consumers. Some of the disclosures that we received from companies operating in this field referred to this decision when they declined to provide detailed information on their algorithms. In the field of consumer contract law, German courts are also not known to be supportive of innovative disclosure designs. Empirical evidence suggests that businesses are currently cautious about employing one-pagers for fear that courts might not approve of them if they were subjected to judicial review.⁹⁷ Under the current legal framework, it is thus only reasonable to expect that businesses would implement the suggestions for improvement reviewed above as non-binding additions to their current privacy policies and terms and conditions, for example in the form of FAQs, explainer videos, or the like.⁹⁸ It remains to be hoped that increasing public awareness and policy makers’ current prioritisation of AI ethics, control, and oversight will motivate data protection authorities and consumer organizations to better monitor and possibly evoke market compliance, thus supporting our calls for better AI

⁹³ EDBP, EDPS (2021).

⁹⁴ Wulf and Seizov (n 7).

⁹⁵ Uber drivers v. Uber B.V. (n 29).

⁹⁶ We were only able to read the unofficial English translation of the judgement.

⁹⁷ Wulf and Seizov (n 22).

⁹⁸ Seizov and Wulf (n 42).

transparency. There is surely work to be done in this field: Facebook, the world's largest social network, never even bothered to reply to our information requests regarding their use of AI algorithms, thereby ignoring its obligations under the GDPR.

Acknowledgements We sincerely thank Shoko Suzuki, Sönke Häselser and Rolf Schwartmann for their valuable comments and suggestions. All errors are the authors' responsibility.

Funding Open Access funding enabled and organized by Projekt DEAL. This research was funded by the DFG (German Research Foundation; project number: WU 824/1-1). It is part of the research project "The ABC of Online Disclosure Duties. Towards a More Uniform Assessment of the Transparency of Consumer Information in Europe" funded by the DFG and the NWO. We thank Joasia Luzak, Marco Loos and Mia Junuzović for their cooperation in this research project.

Availability of data and material Survey and content analysis data are available.

Declarations

Conflict of interest The authors declare that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adadi A, Berrada M (2018) Peeking inside the black-box: a survey on Explainable Artificial Intelligence (XAI). *IEEE Access* 6:52138–52160
- Arakawa Y (2017) Empirical research on human behaviour change and digital intervention through maintaining one-way car-sharing. *Int J Serv Knowl Manag* 1:31–42
- Balasubramanian R, Libarikian A, McElhaney D (2018) Insurance 2030—The impact of AI on the future of insurance. McKinsey & Company
- Bateman J, Wildfeuer J, Hiippala T (2017) Multimodality: Foundations, research and analysis – A problem-oriented introduction. Walter de Gruyter, Berlin
- Ben-Shahar O, Chilton A (2016) Simplification of privacy disclosures: an experimental test. *J Leg Stud* 45:S41–S67
- Ben-Shahar O, Schneider CE (2014) More than you wanted to know: The Failure of Mandated Disclosure. Princeton University Press, Princeton
- Berger-Walliser G, Barton TD, Haapio H (2017) From visualization to legal design: a collaborative and creative process. *Am Bus LJ* 54:347–392
- Bhatt U (2018) Maintaining The Humanity of Our Models. In: 2018 AAAI Spring Symposium Series
- Buyers J (2018) Artificial intelligence: the practical legal issues. Law Brief Publishing, Minehead
- Camilleri MA (2018) Market segmentation, targeting and positioning. *Travel marketing, tourism economics and the airline product*. Springer, New York, pp 69–83
- Chromik M, Eiband M, Völkel ST, Buschek D (2019) Dark Patterns of Explainability, Transparency, and User Control for Intelligent Systems. In: *IUI Workshops, 2019*
- Cooper S (2018) Data science for beginners. CreateSpace, Manchester
- Dittmar T, Hilbert A (2015) Bonitätsprüfung mit Hilfe Künstlicher Neuronaler Netze. *Zeitschrift für Bankrecht und Bankwirtschaft* 10:343–352
- Dorffmeister L (2017) Die europäische Wohnimmobilienkreditrichtlinie. *Ifo Schnelldienst* 70:41–44
- EDPB, EDPS (2021) EDPB-EDPS Joint Opinion 5/2021 on the proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). EDPB/EDPS, Brussels
- Elshout M, Elsen M, Leenheer J, Loos M, Luzak J (2016) Study on Consumers' Attitudes Towards Terms and Conditions (T&Cs). European Commission, Brussels
- Enzi M, Parragh SN, Pisinger D, Prandstetter M (2020) Modeling and solving the multimodal car-and ride-sharing problem. arXiv preprint arXiv:200105490
- European Commission (2019) Guidance on the Interpretation and Application of Council Directive 93/13/EEC on Unfair Terms in Consumer Contracts. European Commission, Brussels
- Eurostat (2021) E-commerce statistics for individuals. Eurostat, Luxembourg
- Faust F, Schäfer HB (2019) Zivilrechtliche und rechtsökonomische Probleme des Internet und der künstlichen Intelligenz. Mohr Siebeck, Tübingen
- Feiler L, Forgó N, Weigl M (2018) The EU General Data Protection Regulation (GDPR): A Commentary. Globe Law and Business, Woking
- Finlay S (2017) Artificial intelligence and machine learning for business: A no-nonsense guide to data-driven technologies. Relativistic Books, London
- Floridi L et al (2018) AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Mind Mach* 28:689–707
- Forbes (2020) Forbes Global 2000, <https://www.forbes.com/global/2000/#103b36be335d>
- Furnell S, Phippen A (2012) Online privacy: a matter of policy? *Comput Fraud Secur* 2012:12–18
- Gellert R et al (2021) The Ola & Uber judgments: for the first time a court recognises a GDPR right to an explanation for algorithmic decision-making. *EU Law Analysis*, <https://eulawanalysis.blogspot.com/2021/04/the-ola-uber-judgments-for-first-time.html>
- Gierschmann S, Schlender K, Stentzel R, Veil W, Gaitzsch P, Buchholtz G, Moser J (2017) Kommentar Datenschutz-Grundverordnung (E-Book). Bundesanzeiger Verlag, Köln
- Helbing D et al (2019) Will democracy survive big data and artificial intelligence? Towards digital enlightenment. Springer, New York, pp 73–98
- High-Level Expert Group on AI (2019) Ethics Guidelines for Trustworthy AI. European Commission, Brussels
- Kulmann F, Reucher E (2000) Computergestützte Bonitätsprüfung bei Banken und Handel. *DBW Die Betriebswirtschaft* 60:113–122
- Laurer M, Seidl T (2021) Regulating the European data-driven economy: a case study on the general data protection regulation. *Policy & Internet* 13(2):257–277

- Luzak JA (2014) Privacy notice for dummies? Towards European guidelines on how to give “clear and comprehensive information” on the cookies’ use in order to protect the internet users’ right to online privacy. *J Consum Policy* 37:547–559
- Marotta-Wurgler F (2011) Will Increased Disclosure Help? Evaluating the Recommendations of the ALI’s Principles of the Law of Software Contracts. *U Chi L Rev* 78:165–186
- Maslen J (2019) Car-sharing giants car2go and DriveNow rebrand as ShareNow as owners BMW and Daimler invest €1bn in mobility services merger. *Asset Finance International*. <https://www.assetfinanceinternational.com/index.php/fleet-finance/fleet-emea/fleet-emea-articles/18005-car-sharing-giants-car2go-and-drive-now-rebrand-as-share-now-as-owners-bmw-and-daimler-invest-1bn-in-mobility-services-merger>
- Mengelkamp AJ (2017) Informationen zur Bonitätsprüfung auf Basis von Daten aus sozialen Medien. Cuvillier Verlag, Göttingen
- OECD AI Policy Observatory (2019) OECD Principles on AI. OECD, Paris
- OECD (2020) The impact of Big Data and Artificial Intelligence (AI) in the insurance sector. OECD, Paris <https://www.oecd.org/finance/impact-big-data-ai-in-the-insurance-sector.htm>
- Olhede SC, Wolfe PJ (2018) The growing ubiquity of algorithms in society: implications, impacts and innovations. *Philos Trans R Soc Math Phys Eng Sci* 376:20170364
- Pollach I (2005) A typology of communicative strategies in online privacy policies: ethics, power and informed consent. *J Bus Ethics* 62:221–235
- Rossov A (2018) The birth of the GDPR: What is it and what you need to know. *Forbes*. <https://www.forbes.com/sites/andrewrossow/2018/05/25/the-birth-of-gdpr-what-is-it-and-what-you-need-to-know/#2024c9a755e5>
- Schwartmann R, Hermann M, Mühlenbeck RL (2020) Transparenz bei Medienintermediären. Vistas, Leipzig
- Seizov O, Wildfeuer J (2017) New studies in multimodality: conceptual and methodological elaborations. Bloomsbury Academic, London, New York
- Seizov O, Wulf AJ (2020) Communicating legal information to online customers transparently: a multidisciplinary multistakeholderist perspective. *J Int Consum Mark* 33:155–179
- Seizov O, Wulf AJ, Luzak J (2019) The transparent trap: a multidisciplinary perspective on the design of transparent online disclosures in the EU. *J Consum Policy* 42:149–173
- Shahriari K, Shahriari M (2017) IEEE standard review—Ethically aligned design: a vision for prioritizing human wellbeing with artificial intelligence and autonomous systems. In: 2017 IEEE Canada International Humanitarian Technology Conference (IHTC), 2017. IEEE, pp 197–201
- Statista (2020a) Größte deutsche Unternehmen nach ihrem weltweiten Umsatz im Geschäftsjahr 2019/2020a, <https://de.statista.com/statistik/daten/studie/12917/umfrage/rangliste-der-500-groessten-unternehmen-deutschlands>
- Statista (2020b) Market share of the biggest Internet companies worldwide in June 2019. <https://de.statista.com/statistik/daten/studie/217485/umfrage/marktwert-der-groessten-internet-firmen-weltweit>
- Suzuki S (2018) Technological civilization and human society in the AI era - AI technology and human future. *Journal of Information and Communication Policy* 2
- Thomas (2020) 25 AI insurance companies you should know about, *BuiltIn*, <https://builtin.com/artificial-intelligence/ai-insurance>
- UNESCO (2019) Steering AI and Advanced ICTs for Knowledge Societies. UNESCO, Paris
- van Boom WH, van der Rest J-PI, van den Bos K, Dechesne M (2020) Consumers beware: online personalized pricing in action! how the framing of a mandated discriminatory pricing disclosure influences intention to purchase. *Soc Justice Res* 33:331–351
- Van Dijck J (2014) Datafication, dataism and dataveillance: big data between scientific paradigm and ideology. *Surveill Soc* 12:197–208
- Veale M, Edwards L (2018) Clarity, surprises, and further questions in the Article 29 Working Party draft guidance on automated decision-making and profiling. *Comput Law Secur Rev* 34:398–404
- Wachter S, Mittelstadt B, Floridi L (2017) Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *Int Data Privacy Law* 7:76–99
- Waller R (2017) Graphic literacies for a digital age. In: *Information Design*. Routledge, London, pp 193–220
- Whittlestone J, Nyrop R, Alexandrova A, Dihal K, Cave S (2019) Ethical and societal implications of algorithms, data, and artificial intelligence: a roadmap for research. Nuffield Foundation, London
- Wulf AJ (2014) Institutional competition of optional codes in European contract law. *Eur J Law Econ* 38:139–162
- Wulf AJ (2016) The contribution of empirical research to law. *J Jurisprudence* 29:29–49
- Wulf AJ, Seizov O (2020a) Artificial intelligence and transparency: a blueprint for improving the regulation of AI applications in the EU. *Eur Bus Law Rev* 31:611–640
- Wulf AJ, Seizov O (2020b) The principle of transparency in practice. How different groups of stakeholders view EU online information obligations. *Eur Rev of Private Law* 20:1065–1092

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.