WRITTEN BY CHRISTINA HUCHRO

Article36

WWW.ARTICLE36.ORG
INFO@ARTICLE36.ORG
@ARTICLE36

Article 36 is a specialist non-profit organisation, focused on reducing harm from weapons.

FOUNDATIONS OF THE GDPR: PRINCIPLES RELEVANT TO DISCUSSIONS ON AUTONOMOUS WEAPONS

KEY MESSAGES

- During the development of the General Data Protection Regulation (GDPR), European legislators recognized the need to protect individuals from harms resulting from automated decision-making.
- Legislators included a general prohibition (subject to limited exceptions) on "decisions based solely on automated processing" in order to: 1) prevent data controllers from making legally significant decisions solely on the basis of an individual's "data shadow;" and 2) avoid situations where data controllers abdicate their responsibility by placing too much trust in outcomes reached through automated processes. The prohibition essentially establishes a presumption that decisions based solely on automated processes are unacceptable, and that controllers bear the burden of demonstrating compliance.
- States that have adopted the GDPR have thus recognized the foundational argument that outcomes reached solely through automated processes threaten to violate human dignity by upending an individual's right to "constitute" themselves.
- As states have sought to mitigate these risks in the civil context through the passage of the GDPR, legislative steps should be taken to address the very similar concerns which arise in the context of armed conflict regarding the use of automated weapons systems (AWS).

I. THE GDPR AND THE ESTABLISHMENT OF ALGORITHMIC ACCOUNTABILITY

The European Union's adoption of the General Data Protection Regulation (GDPR) in April 2016 marked the culmination of a decades-long effort to establish new comprehensive rules on the processing of personal data.[1] The regulation, which came into effect in May 2018, quickly became known as "the most consequential regulatory development in information policy in a generation,"[2] and has had a dramatic influence on how countries throughout the world approach data protection.[3]

Of most relevance to policy debates surrounding autonomous weapons systems (AWS), the GDPR includes a series of new regulations which address algorithmic accountability.[4] As scholars have noted, there are numerous risks associated with unchecked automated processes; for example, automated decision-making, "can often be opaque, complex, and subject to error, bias, discrimination, in addition to implicating dignitary concerns."[5] The GDPR seeks to address these concerns by: 1) imposing obligations on "data controllers" who create risks of harm through automated decision-making; and 2) establishing new rights for "data subjects" whose interests may be affected by the controllers' automated processing. These obligations and rights are captured by Articles 13, 14, 15, 22, and 35, each of which refers specifically to "automated decision-making" or "automated processing."

DEFINING 'AUTOMATED DECISION-MAKING'

Before addressing the scope of these regulations, it is necessary to first clarify how the GDPR defines "automated decision-making." As the aim of this paper is to establish parallels between the concerns which motivated legislators to prohibit automated decision-making in the data protection context, and similar concerns about automation in the AWS space, it is important to clearly define what "automated" means in both contexts.

The Article 29 Data Protection Working Party,[6] an advisory body now known as the European Data Protection Board (EDPB), provided guidelines on what constitutes a "decision based solely on automated processing.[7]" In the *'Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679,'* the Working Party explained:

> "An automated process produces what is in effect a recommendation concerning a data subject. If a human being reviews and takes account of other factors in making the final decision, that decision would not be 'based solely' on automated processing . . . To qualify as human involvement, the controller must ensure that any oversight of the decision is meaningful, rather than just a token gesture. It should be carried out by someone who has the authority and competence to change the decision.[8]"

In simpler terms, European legislators were concerned about situations where data controllers automatically accept recommendations generated through a machine process, without imposing additional human checks or safeguards. Thus, legislators established a general prohibition on automated decision-making, allowing them only where there is meaningful human control.

Scholars in the AWS space approach the term "automated" in a very similar way, except, rather than focusing primarily on the need for human review *after* a decision has been reached, scholars focus on the degree of control or input-limitations needed *before* force is applied. The starting point of concerns is framed in a similar way: weapons in which, "force will be applied on the basis of data collected by sensors, without human evaluation of that data, and without a human setting the time and place of that application of force."[9] Control measures are then proposed in relation to the duration and spatial area of that process, and in terms of how things might be identified as targets. This approach—allowing for certain forms of automated decision-making where human intervention at the outset limits the range of possibilities—is essentially the *mirrored version* of the approach taken by the GDPR, which generally prohibits automated decision-making unless there is an opportunity for human intervention *after* the fact. An exception to this might be for systems that would *target people directly*, where the policy approaches of Article 36, Stop Killer Robots and the ICRC amongst others suggest a blanket prohibition on automated targetting more in line with the GDPR approach.[10]

The approach taken in the GDPR is arguably more restrictive than that adopted in the policy discussion on autonomous weapons. However, both of these approaches highlight legislative fears about pure automation and the absence of human intervention; thus, it follows that the concerns which motivated legislators to prohibit certain forms of automated decision-making in the data protection context, are applicable to the AWS space.

DEFINING 'PROFILING'

While not the exact purview of this paper, it is also important to define how the GDPR approaches "profiling," as many of the provisions on automated decision-making in the GDPR simultaneously address profiling issues. As defined by Article 4 of the GDPR, profiling involves:

> "Any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyze or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements."[11]

The Article 29 Working Party has explained that profiling involves three distinct elements: 1) an automated form of processing; 2) the use of personal data; and 3) an objective to evaluate personal aspects about a person.[12] In regard to the third criteria, the Working Party noted that a data controller must have the *purpose* of using data on an individual's characteristics, "to place them into a certain category or group, in particular to analyze and/or make predictions about."[13] Consequently, if a controller merely collects data about a user's race, gender, or religion to gain a general composition of their clientele, it would not be profiling because the controller is not using the data to make predictions about certain behavior or to assess likely conduct.[14]

While "profiling" and "automated decision-making" are often discussed together in the context of the GDPR, the scope of these two activities can differ because "profiling does not consist of a *decision* but rather an *evaluation* of personal aspects about an individual.[15]" Profiling is essentially a decision-making tool used to assist or to inform data controllers about certain decisions. However, it is important to note that automated decision-making can occur *without* profiling. For example, "[i]mposing speeding fines purely on the basis of evidence from speed cameras is an automated decision-making process that does not necessarily involve profiling."[16]

Although the concerns underlying the provisions on automated decision-making are more directly relevant to the AWS debate, and the legislative emphasis on profiling waned throughout the drafting process, the GDPR's treatment of profiling has significant parallels to autonomous weapons.[17] Specifically, the GDPR focuses on controllers who use data about an individual's characteristics to predict how that subject will behave. Scholars in the AWS space express similar concerns about "target profiles." Target profiles consist of "a set of conditions which result in a specific application of force being undertaken by the system.[18]" For example, if on-the-ground sensors collect information which meets the pre-defined "target profile," the system can apply force. In essence, these systems are making predictions about who these individuals are, and how they are likely to behave, based on patterns of data on the ground; the system could theoretically predict that a certain individual is an enemy combatant, and apply force against that individual, because the data collected on them matches the target profile of an enemy combatant. Thus, the practice of using target profiles in automated weapons is very similar to how controllers use profiling in the data protection context: both use a necessarily incomplete set of data to make predictions about who an individual is (or what an object is), leading to potentially life-altering consequences. While this paper will focus more on the automation provisions of the GDPR, the provisions on profiling are clearly applicable to ongoing discussions in the AWS policy space.

PROVISIONS ON AUTOMATED DECISION-MAKING AND PROFILING

Articles 13, 14, 15, 22, and 35 of the GDPR each address a certain aspect of automated decision-making. Articles 13, 14, and 15 function primarily to introduce more transparency into automated processes. Article 35 requires data controllers to conduct *ex ante* risk assessments, in the form of a Data Protection Impact Assessment (DPIA), prior to

utilizing new forms of processes that are likely "to result in a high risk to the rights and freedoms of natural persons."[19] Finally, while subject to significant scholarly debate, Article 22 has been interpreted as a general prohibition on certain forms of automated decision-making. Together, these regulations provide significantly more protection for data subjects who may experience harm as a result of automated processes.

ARTICLES 13 AND 14

Articles 13 and 14 enhance transparency within data processing by requiring that data controllers "ensure they explain clearly and simply to individuals how the profiling or automated decision-making process works."[20] Specifically, Articles 13(2)(f) and 14(2)(g) impose notification obligations on controllers. In order to "ensure fair and transparent processing," data controllers must disclose: "the *existence* of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, *meaningful information* about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject."[21]

There has been significant scholarly debate on what data controllers need to provide to satisfy the "meaningful information" threshold, especially in the context of algorithms which may be challenging to explain to the general public. As the Article 29 Working Party Guidelines note, however, data controllers do not need to provide "a complex explanation of the algorithms used or disclosure of the full algorithm."[22] Instead, controllers must provide information which is "sufficiently comprehensive for the data subject to understand the reasons for the decision."[23] Scholars have argued that the phrase "meaningful information" should be interpreted *flexibly*, such that data subjects are able to exercise their rights without restricting the ability of data controllers to use innovative forms of artificial intelligence and machine learning.[24] A flexible interpretation is also necessary given the fact that Member States have translated the phrase "meaningful information" differently. For example, "[t]he German text of the GDPR uses the word *'aussagekräftige [informationen]'*, the French text refers to *'informations utiles'*, and the Dutch version uses *'nuttige informative'*. . . [t]hese are related, but not identical concepts, suggesting that a flexible, functional approach will be most appropriate."[25] Specifically, while the German words translate directly into the phrase "*meaningful* information," the French and Dutch formulations are translated in English to mean "*useful* information."[26]

ARTICLE 15

In contrast to Articles 13 and 14, Article 15 is linguistically structured as a *right* for data subjects, rather than an *obligation* imposed on data controllers.[27] Specifically, data subjects have the "right to obtain from the controller" information regarding "the existence of automated-decision-making."[28] Scholars have noted that the difference between Articles 13, 14, and 15 is that, "Articles 13 and 14 might require an overview of a system prior to processing, but Article 15's access right could provide *deeper disclosure*, including insight into a particular decision affecting a particular individual."[29]

ARTICLE 35

Article 35 requires data controllers to conduct a Data Protection Impact Assessment (DPIA), "where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons."[30] In short, data controllers have an obligation to assess the risks posed by automated decisions and determine "whether the processing operations are necessary and proportionate, and the remedial measures planned to deal with the risks."[31] Of particular relevance to the AWS debate, GDPR legislators recognized that requiring data controllers to clearly explain the purpose, necessity, safeguards, and impacts of automated processing introduces more accountability into the process.[32] In simpler terms, the inclusion of DPIA requirements, "makes it clear that the legislator ascribes a *high risk* to automated decisions."[33]

ARTICLE 22

Finally, the provision which is most relevant to the AWS debate is Article 22. Article 22 states that data subjects "shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her."[34] There are three exceptions to this requirement. However, even if one of the exceptions apply, the data controller must implement measures to safeguard the data subject's rights.[35]

Since the GDPR's passage, there has been significant scholarly debate on whether Article 22 should be interpreted as "a right to object to decisions [reached through automated means] or a general prohibition on significant algorithmic decision-making."[36] On this issue, the Article 29 Data Protection Working Party guidelines provide significant clarification:

> "Interpreting Article 22 as a *prohibition* rather than a *right* to be invoked means that individuals are automatically protected from the potential effects this type of processing may have. The wording of the Article suggests that this is the intention and is supported by Recital 71."[37]

While the interpretive guidelines from the Article 29 Working Party do not have the "direct force of law . . . [t]hey are, nonetheless, strongly indicative of how enforcers will interpret the law."[38]

II. LEGISLATIVE CONCERNS RELATED TO AUTOMATED DECISION-MAKING

As articulated above, the GDPR has successfully implemented a range of tools aimed at curtailing the potentially harmful effects of automated decision-making; these Articles, and in particular Article 22, have already had a significant impact on how data controllers use automated forms of processing. The concerns that motivated the European Commission to establish such robust safeguards, which are highly relevant and analogous to concerns in the AWS space, are discussed in further detail below.

HISTORICAL ORIGINS OF ARTICLE 22 OF THE GDPR

The origins of Article 22 of the GDPR can be traced back to France's 1978 Act on Data Processing, Files and Individual Liberties *(Loi no.*

78-17 du 6. Janvier 1978),[39] which stated that "information technology should be at the service of every citizen. Its development shall take place in the context of international co-operation. It shall not violate human identity, human rights, privacy, or individual or public liberties. This Act shall apply to *automatic processing* of personal data as well as non-automatic processing of personal data."[40]

France's pioneering 1978 Act influenced the development of Article 15 of the 1995 EC Directive on Data Protection (Directive 9 5/46 or "DPD"). Article 15 of the DPD is the predecessor to Article 22 of the GDPR.[41] Indeed, the first draft of the GDPR,[42] which was tabled by the European Commission on January 15, 2012, stated, "Article 20 [now Article 22] concerns the data subject's right not to be subject to a measure based on profiling. It builds on, with modifications and additional safeguards, Article 15(1) of Directive 9 5/46 on automated individual decisions."[43]

Article 15 of the DPD, which was the "first pan-European legislative norm aimed directly at regulating *purely machine-based decisions* in a data protection context,[44]" states:

> "Member States shall grant the right to every person not to be subject to a decision which produces legal effects concerning him or significantly affects him and which is based solely on automated processing of data intended to evaluate certain personal aspects relating to him, such as his performance at work, creditworthiness, reliability, conduct, etc."[45]

The concerns that led to the passage of Article 15 of the DPD, the predecessor to Article 22 of the GDPR, were first articulated in the *Council Directive Concerning the Protection of Individuals*[46] and the *Proposal for a Council Directive Concerning the Protection of Individuals in relation to Processing of Personal Data.*[47] These two documents are particularly important because there are very few recent documents which reveal the exact legislative rationale behind Article 22 of the GDPR.[48] However, as scholars Isak Mendoza and Dr. Lee A. Bygrave noted, "it is safe to assume that the rationale [behind Article 22] is at least partly rooted in the concerns that gave rise to DPD Art. 15 over two decades ago."[49]

Two primary concerns emerge within these documents.[50] First, the drafters were concerned that the proliferation of automated decision-making would "diminish the role played by persons in shaping important decision-making processes that affect them."[51] Second, drafters feared that "increasing automatization of decision-making processes engenders automatic acceptance of the validity of the decisions reached and a concomitant reduction in the investigatory and decisional responsibilities of humans."[52] The following sections address these concerns in tandem and highlight relevant connections to the AWS debate.

CONNECTIONS BETWEEN THE CONCERNS MOTIVATING THE PASSAGE OF ARTICLE 22 AND AUTONOMOUS WEAPONS

CONCERN 1: THREATS POSED BY 'DATA SHADOWS'

Summary: 'Data shadows' can be erroneous and are inherently incomplete, and decisions based entirely on 'data shadows' threaten to undermine human dignity by upending an individual's right to 'constitute' themselves.

Background: In the *'Proposal for a Council Directive Concerning the Protection of Individuals in relation to the Processing of Personal Data,'* European legislators emphasized the dangers posed by automated decision-making. Legislators were concerned about the "potential weakening of the ability of persons to exercise influence over decision-making processes that significantly affect them, in light of the growth of automated profiling practices.[53]" The legislators note:

> "This provision is designed to protect the interest of the data subject in participating in the making of decisions which are of importance to him. The use of extensive data profiles of individuals by powerful public and private institutions deprives the individual of the capacity to influence decision-making processes within those institutions, should decisions be taken on the sole basis of his 'data shadow.'"[54]

Several scholars have connected legislative concerns about decisions taken solely on the basis of a "data shadow" with more foundational concerns for human dignity.[55] For example, in analyzing the Commission's explanation, Mendoza and Bygrave commented, "reading between the lines of these explanatory statements, we can discern not just fear about humans letting machines make mistakes but a concern to uphold human dignity by ensuring that humans (and not their 'data shadows') maintain the primary role in 'constituting' themselves."[56] In a separate paper, Bygrave posits the Commission was concerned that, "the registered data images of persons (their "data shadows") *threaten to usurp the constitutive authority of the physical self* despite their relatively attenuated and often misleading nature."[57] In other words, granting data controllers the ability to make legally significant decisions solely on the basis of 'data shadows' undermines human dignity by subjecting individuals to harms based on potentially misleading, incomplete, or inaccurate information.[58]

Connection to Autonomous Weapons Systems: Concerns about the risks posed by decisions triggered solely on the basis of a "data shadow" are highly relevant to the AWS debate. For background purposes, a "data shadow" consists of the "collective body of data that is automatically generated and recorded [by] . . . sensors and IP surveillance, metadata from communications and security and authentication mechanisms."[59] In formulating Article 15 of the DPD, and presumably Article 22 of the GDPR, legislators wanted to prevent data controllers from making decisions *solely* on the basis of this collection of data. Legislators recognized that "mechanical determinations" derived from data shadows can often be erroneous.[60] Scholars have articulated similar concerns when discussing autonomous weapons, especially weapons in which, "the decision to attack is implemented through *sensors* that acquire data in the world, algorithms that process and classify sensor data according to pre-encoded profiles of targets, and actuators that apply kinetic or other force to targets."[61] In essence, these systems have the potential to make life-and-death decisions based on "data shadows," or data collected through on-the-ground sensors.

There are subtle temporal differences between "data shadows" formed in the data protection context and "data shadows" formed in the AWS space. In the digital world, data shadows are developed over an extended period of time and have the potential to "linger" for decades.[62] In contrast, "data shadows" in the AWS space are often collected and formulated based on real-time data on the ground; thus, decisions are made on a much more imminent, short-term basis. However, the

temporal differences between these two formulations do not alter the fact that the concerns about data shadows in the data protection context are highly relevant to negotiations in the AWS space. Specifically, data shadows can be erroneous, are necessarily *incomplete* representations of real humans, and can deprive individuals of their fundamental right to constitute themselves.

CONCERN 2: MISPLACED TRUST IN MACHINE-GENERATED OUTCOMES

Summary: Over-reliance on machine-generated outcomes can negate human responsibility and accountability, undermining the ability of data subjects to seek appropriate legal redress.

Background: In passing Article 15 of the DPD, legislators expressed concerns about the role of human accountability and responsibility during automated processing. Specifically, in the explanatory text for the *Proposal for a Council Directive Concerning the Protection of Individuals in relation to the Processing of Personal Data*, the Commission explained:

> "The danger of the misuse of data processing in decision-making may become a major problem in future: the result produced by the machine, using more and more sophisticated software, and even expert systems, has an apparently objective and incontrovertible character to which a human decision-maker may attach too much weight, thus abdicating his own responsibilities. Article 16 (1) therefore lays down the principle that a person is not obliged to accept a decision of a public administration or of a private party which adversely affects him if it is based solely on automatic processing defining a personality profile."[63]

Legislators sought to combat the false narrative that "machines could not display the biases of people and so would be ideal neutral decision makers."[64] The final version of Article 22 of the GDPR directly responds to these concerns by requiring that humans are ultimately responsible for safeguarding the rights of data subjects. In explaining what constitutes appropriate "safeguards" under Articles 22(2)(b), 22(3), and 22(4), the Article 29 Working Party states:

> "[h]uman intervention is a *key element*. Any review must be carried out by someone who has the appropriate authority and capability to change the decision. The reviewer should undertake a thorough assessment of all the relevant data, including any additional information provided by the data subject."[65]

In short, the text of the GDPR directly mandates that humans ultimately remain *responsible* for the potential harms created by automated processing; this, in turn, ensures data subjects who have been harmed can seek appropriate legal remedies.

Connection to Autonomous Weapons Systems: Concerns regarding both the over-reliance on machine-generated outcomes and the negation of human responsibility are similarly relevant to the AWS debate. For example, in a report on the risks posed by data issues in military autonomous systems, the United Nations Institute for Disarmament Research (UNIDIR) detailed the threats posed by reliance on automated processes; the report notes, "[a] badly calibrated or faulty sensor feeding an autonomous system might generate an incorrect measurement (such as the size, shape or speed of a target), or a human-generated data feed may include errors (incorrect numbers, spelling mistakes, incorrect formatting, etc.)."[66] In other words, results generated through machine processes cannot be viewed as "objective" or "incontrovertible" given the potential risks posed by incomplete or incorrect data.

Regarding the negation of human responsibility and legal redress, autonomous weapons raise serious issues related to attribution and compliance with international humanitarian law (IHL). Autonomous weapons can be unpredictable and obscure, both of which pose challenges for legal compliance. For example, as the International Committee of the Red Cross (ICRC) explained, "if an AWS' functioning is opaque, then humans responsible for the application of IHL rules – both persons entrusted with the legal review of an AWS and persons responsible for compliance with IHL during its use – could not reasonably determine its lawfulness under IHL."[67] In other words, there is an "accountability gap" whereby AWS users could potentially be held legally or administratively responsible for unpredictable outcomes generated by autonomous weapons, even though personally they lacked the requisite intent and understanding of how these systems work.[68] Thus, in both the data protection context and the AWS space, requiring meaningful human review ensures both that humans remain responsible and accountable for outcomes generated through automated processes and that there is alignment between legal redress and personal culpability.

III. CONCLUSION

The concerns which motivated legislators to implement new regulations for automated decision-making in the data protection context are highly relevant and applicable to ongoing policy debates in the AWS space. By accepting and implementing the GDPR's provisions on algorithmic accountability, European states have recognized that subjecting individuals to automated decision-making, without meaningful human oversight, may violate human dignity. These concerns emerge most prominently in Article 22, which establishes a presumption that certain decisions based purely on automatic processes are unacceptable, and Article 35, which places an obligation on the data controller to ensure that automated processes adequately safeguard the rights of data subjects. In short, states recognized foundational concerns about automation in the civil context and responded by establishing a presumption that such systems are unacceptable unless proven otherwise by the controller; the controller bears the burden of demonstrating compliance. Similar attention, and perhaps a similar legislative approach whereby autonomous weapons are presumed unacceptable unless certain conditions can be met, should be further considered in the AWS space.

ENDNOTES

1 *The History of the General Data Protection Regulation*, European Data Protection Supervisor (last visited Jan. 18, 2022), https://edps.europa.eu/data-protection/data-protection/legislation/history-general-data-protection-regulation_en.

2 Chris Jay Hoofnagle, Bart van der Sloot & Frederik Zuiderveen Borgesius, *The European Union General Data Protection Regulation: What it is and What It Means*, 28 Information & Communications Technology Law 65, 66 (2019).

3 *See, e.g.*, Oskar Josef Gstrein, *The Internet on its Way Back to a Future of Human Dignity?*, Völkerrechtsblog (May 29, 2019), https://voelkerrechtsblog.org/the-internet-on-its-way-back-to-a-future-of-human-dignity/ ("Despite the fact that Europe plays a second-tier role when it comes to the development and production of data-driven technology and services, research shows that of 132 countries which have national data protection laws at the beginning of 2019 the majority adopts the European 'omnibus approach'").

4 *See* Regulation (EU) 2016/679, of the European Parliament and the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), 2016 O.J. (L 119) 1 at arts. 13, 14, 15, 22, and 35 [hereinafter GDPR].

5 Margot E. Kaminski, *The Right to Explanation, Explained*, 34 Berkerly Tech. L.J. 189, 191 (2019).

6 First created in 1996 by Article 29 of the Data Protection Directive, the "Article 29 Working Party" is an advisory body composed of national Data Protection Authorities from the Member States, the European Data Protection Supervisor (EDPS), and a representative of the European Commission. For more information, *see* European Commission MEMO/10/542, Data Protection Reform – Frequently Asked Questions (November 4, 2010).

7 For a more detailed account of the responsibilities and role of the Article 29 Working Party, *see* Kaminski, *supra* note 5, at 194 ("Article 29 Working Party guidelines . . . do not have the direct force of law. They are, nonetheless, strongly indicative of how enforcers will interpret the law. Now that the GDPR is in effect, these guidelines have additional, though indirect, teeth. The European Data Protection Board under the GDPR has additional supervisory and harmonizing capabilities over Member State Data Protection Authorities. A local Data Protection Authority, in other words, is now even more likely to adhere to the guidelines than under the Directive").

8 Article 29 Data Protection Working Party, *Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679*, WP 251, 19 (Oct. 3, 2017, revised Feb. 6, 2018) [hereinafter Article 29 Guidelines on Automated Decisions].

9 Anna Turek & Richard Moyes, *Autonomy in Weapons: 'Explicability' as a Way to Secure Accountability*, Article 36 (Dec. 2020), https://article36.org/wp-content/uploads/2020/12/Explicability-and-accountability.pdf.

10 *See e.g.* Maya Brehm, *Targeting People: Key Issues in the Regulation of Autonomous Weapons Systems*, Article 36 (Nov. 2019), https://article36.org/wp-content/uploads/2019/11/targeting-people.pdf, Stop Killer Robots, *Our Policy Position*, https://www.stopkillerrobots.org/our-policies/, ICRC Position on Autonomous Weapons Systems, ICRC (May 12, 2021), https://www.icrc.org/en/document/icrc-position-autonomous-weapon-systems

11 GDPR at art. 4(4).

12 Article 29 Guidelines on Automated Decisions, *supra* note 8, at 6–7.

13 *Id.* at 7.

14 For additional examples on what constitutes 'profiling,' *see* Article 29 Guidelines on Automated Decisions, *supra* note 8, at 8 ("A data broker collects data from different public and private sources, either on behalf of its clients or for its own purposes. The data broker compiles the data to develop profiles on the individuals and places them into segments. It sells this information to companies who wish to improve the targeting of their goods and services. The data broker *carries out profiling* by placing a person into a certain category according to their interests.").

15 Kim Lucassen, *Rights Related to Automated Decision-Making and Profiling Under the GDPR and D-DPA*, Loyens & Loeff (Aug. 13, 2018), https://www.loyensloeff.com/en/en/news/news-articles/rights-related-to-automated-decision-making-and-profiling-under-the-gdpr-and-d-dpa-n10518/.

16 Article 29 Guidelines on Automated Decisions, supra note 8, at 6–8; *see also Data Is Power: Profiling and Automated Decision-Making in GDPR*, Privacy International (2017), https://privacyinternational.org/sites/default/files/2018-04/Data%20Is%20Power-Profiling%20and%20Automated%20Decision-Making%20in%20GDPR.pdf ("Profiling practices create, discover or construct knowledge from large sets of data from a variety of sources. Such knowledge can be used to make or inform decisions that *may or may not be automated*.").

17 The decision to focus primarily on automated decision-making, rather than profiling, is supported by the legislative history of the GDPR. While early versions of the GDPR focused explicitly profiling, the final version references 'profiling' as an *example* of a form of automated processing. For a more detailed account of this shift, *see* Maja Brkan, *Do Algorithms Rule the World? Algorithmic Decision-Making in the Framework of the GDPR and Beyond*, Technology Policy Institute (2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3124901 ("[I]t is interesting to observe how Article 22 GDPR developed throughout the legislative procedure leading to the adoption of the GDPR as it shows the evolution from focusing specifically on profiling to a more general formulation using a broader notion of automated individual decision-making. Differently from the final GDPR, in the initial Commission's proposal, this article, titled '*Measures based on profiling*', regulated profiling based on automated processing and not generally automated decision-making as the provision in the final GDPR does"); *see also*, Luca Belli, Nicolo Zingales & Yasmin Curzi, Glossary of Platform Law and Policy Terms 279 (Internet Governance Forum, Dec. 2021) ("[T]he explicit mention of profiling could be interpreted simply as *illustrative* of one of the possible risks involved in automated processing").

18 Richard Moyes, *Target Profiles*, Article 36 (Aug. 2019), https://article36.org/wp-content/uploads/2019/08/Target-profiles.pdf.

19 GDPR at art. 35(1).

20 Article 29 Guidelines on Automated Decisions, *supra* note 8, at 16.

21 GDPR at arts. 13(2)(f) and 14(2)(g).

22 Article 29 Guidelines on Automated Decisions, *supra* note 8, at 25.

23 *Id.*

24 Andrew D. Selbst & Julia Powles, *Meaningful Information and the Right to Explanation*, 7 International Data Privacy Law 233, 236 (2017).

25 Id.

26 Simon Chesterman, We, the Robots?: Regulating Artificial Intelligence and the Limits of the Law, 159 (Cambridge University Press, 2021).

27 Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a 'Right to an Explanation' is Probably not the Remedy you are Looking for*, 16 Duke Law and Technology Review 18, 52 (2017).

28 GDPR at art. 15(1)(h).

29 Kaminski, *supra* note 5, at 200.

30 GDPR at art. 35(1).

31 Carsen Orwat, *Risks of Discrimination Through the Use of Algorithms*, Institute for Technology Assessment and Systems Analysis (ITAS) and Karlsruhe Institute of Technology (KIT), 78 (2020), https://www.antidiskriminierungsstelle.de/EN/homepage/_documents/download_diskr_risiken_verwendung_von_algorithmen.pdf?__blob=publicationFile&v=1.

32 For a more detailed explanation on how Article 35 establishes greater accountability, *see* Margot E. Kaminski & Gianclaudio Malgieri, *Algorithmic Impact Assessments Under the GDPR: Producing Multi-Layered Explanations*, 11 International Data Privacy Law 125, 131 (Dec. 6, 2020) ("[I]n the context of the GDPR's algorithmic governance regime, the DPIA should be understood as a nexus between the GPDR's two approaches to governing algorithmic decision-making. The DPIA links the GDPR's individual rights to its systemic governance of algorithms").

33 Orwat, *supra* note 31, at 78.

34 GDPR at art. 22(1).

35 For a more detailed analysis of the structure of Article 22, see Céline Castets-Renard, *Accountability of Algorithms in the GDPR and Beyond: A European Legal Framework on Automated Decision-Making*, 30 Fordham Intell. Prop. Media & Ent. L.J 91, 114 (2019) ("In comparison to Article 15 of the Directive 95/46/EC, Article 22, Section 3 of the GDPR sets forth new guarantees. When the exceptions apply,

'the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests.' These rights are a non-exhaustive list of 'suitable measures.' The controller has to respect, at a minimum, the 'right to obtain human intervention,' the right for the data subject to 'express his or her point of view,' and the right 'to contest the decision.' These requirements could be justified by one of the purposes of the GDPR-to improve the protection based on Article 8 of the EU Charter of Fundamental Rights").

36 Kaminski, *supra* note 5, at 196.

37 Article 29 Guidelines on Automated Decisions, *supra* note 8, at 20–21.

38 Kaminski, *supra* note 5, at 194.

39 *See, e.g.,* Karen Yeung & Martin Lodge, Algorithmic Regulation, 249 (Oxford University Press, 2019) ("[A]lthough enacted in the 1990s, [Article 15's] roots reached back at least to the 1970s. An important source of its inspiration was France's 1978 Act on data processing, files, and individual liberties . . . which prohibited judicial, administrative, or personal decisions involving assessment of human behaviour insofar as these were based solely on automatic data processing which defined the profile or personality of the individual concerned").

40 *Loi n° 78-17 du 6 janvier 1978 relative à l'informatique, aux fichiers et aux libertés*.

41 *See* Brkan, *supra* note 17, at 5 ("[t]his provision continues the legacy of the Data Protection Directive, more precisely its Article 15, according to which the data subject equally had the right not to be subject to a decision producing legal effects or significantly affecting him and which is based solely on automated processing of data. While the wording of the provision did not undergo substantial changes with the adoption of the GDPR, the practical importance of the provision increased with augmented use of automated decision-making.")

42 For a more comprehensive history on the development of the GDPR, *see* Gstrein, *supra* note 3.

43 *Proposal for a Regulation of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data (General Data Protection Regulation*), C7-0025/12, (Jan. 25, 2012).

44 Yeung & Lodge, *supra* note 39, at 249.

45 Directive 95/46/EC of the European Parliament and of the Council of 24.10.1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data (O.J. L 281, 23.11.1995).

46 COM (90) 314 final – SYN 287, p. 29.

47 COM (92) 422 final – SYN 287, p. 26.

48 *See* Isak Mendoza & Lee A. Bygrave, *The Right Not to Be Subject to Automated Decisions Based on Profiling* (University of Oslo Faculty of Law, Rsch. Working Paper No. 2017-20, 2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2964855; ("The *travaux préparatoires* to the GDPR provide little explanation of the rationale for these provisions [Article 22]").

49 *Id.* at 6.

50 For additional background on these concerns, *see* Ira S. Rubinstein, *Big Data: The End of Privacy or a New Beginning?*, International Data Privacy Law, Volume 3, Issue 2, May 2013, Pages 74–87 ("Article 15 seems to have been motivated by the *twin concerns* that automated decision making will diminish the role of persons in influencing decisions affecting them and that such decisions are given too much deference (as if the mere fact that they result from sophisticated computer processes makes them more objective)").

51 Lee A. Bygrave, *Minding the Machine: Article 15 of the EC Data Protection Directive and Automated Profiling,* 17 Computer L. & Security Rep. 17, 18 (2001).

52 *Id.* at 18.

53 Mendoza & Bygrave, *supra* note 48, at 7.

54 COM (90) 314 final – SYN 287, p. 29.

55 *See,* Yeung & Lodge, *supra* note 39, at 250 ("Animating both Article 15 and its French antecedents was fear for the future of *human dignity* in the face of machine determinism. Their rationale was grounded in a concern to ensure that humans maintain ultimate control of, and responsibility for decisional processes that significantly affect other humans, and that they thereby maintain the primary role in 'constituting' themselves"); *see also* Recommendation CM/Rec (2010)13 ("the use of profiling techniques without precautions and specific safeguards could d*amage human dignity* and unjustifiably deprive individuals of access to certain goods or services").

56 Mendoza & Bygrave, *supra* note 48, at 7.

57 Bygrave, supra note 51, at 18.

58 *See, e.g.*, Karoline Krenn, *Bringing Context Back into Privacy Regulation and Beyond*, 21 Economic Sociology 43, 44 (Nov. 2019) ("A strong skepticism towards decision processes based on selected pieces of decontextualized information ('the data shadow') already characterized the European Directive of 1995. The partiality and shallowness of such decisions were considered as dehumanizing and making fully automated decisions was forbidden") (internal citations removed).

59 *Data Shadow*, Tech Target (last visited Jan. 18, 2022), https://whatis.techtarget.com/definition/data-shadow.

60 *See* Paul Nihoul & Pieter Van Cleynenbreugel, The Roles of Innovation in Competition Law Analysis 117–118 (2018) (discussing the risks posed by reliance on determinations based entirely on data shadows).

61 Maya Brehm, *Targeting People: Key Issues in the Regulation of Autonomous Weapons Systems*, Article 36 (Nov. 2019), https://article36.org/wp-content/uploads/2019/11/targeting-people.pdf.

62 *See* David J. Hand, *Aspects of Data Ethics in a Changing World: Where Are We Now?*, 6 Big Data 176, 180 (2018) ("These are the data traces that result from ordinary daily activity, such as using a credit card, a travel card, accessing social media, web searching, e-mails, making phone calls, and even interacting with an electronic ''intelligent personal assistant'' such as Amazon's Alexa. Such traces reveal huge amounts about what people get up to, who they interact with, what their interests are, and even what their beliefs are. . . Unlike real shadows, data shadows may *linger for a long time*–causing potential embarrassment and worse far into the future.").

63 COM (92) 422 final – SYN 287, p. 26.

64 Edwards & Veale, *supra* note 27, at 27; *see also*, Emre Bayamlıoğlu, *Transparency of Automated Decisions in the GDPR: An Attempt for Systemisation*, Tilburg Institute for Law, Technology, and Society (TILT), 8 (2018) ("The human belief in the infallibility of mechanistic processes legitimatizes the biases, and obfuscates the inherent risks").

65 Article 29 Guidelines on Automated Decisions, *supra* note 8, at 27.

66 Arthur Holland Michel, *Known Unknowns: Data Issues and Military Autonomous Systems*, United Nations Institute for Disarmament Research 4 (2021), https://doi.org/10.37559/SecTec/21/AI1.

67 ICRC Position on Autonomous Weapons Systems, ICRC (May 12, 2021), https://www.icrc.org/en/document/icrc-position-autonomous-weapon-systems.

68 *See* Brehm, *supra* note 61.

Acknowledgements:

Christina Huchro is currently a second-year student at Harvard Law School. Prior to attending law school, she worked as an associate at a business advisory firm in Washington, DC. She focuses her research on the intersection between international law, national security, and technology.

chuchro@jd23.law.harvard.edu