



The interrelation between data and AI ethics in the context of impact assessments

Emre Kazim¹ · Adriano Koshiyama¹

Received: 9 July 2020 / Accepted: 2 November 2020 / Published online: 22 November 2020
© The Author(s) 2020

Abstract

In the growing literature on artificial intelligence (AI) impact assessments, the literature on data protection impact assessments is heavily referenced. Given the relative maturity of the data protection debate and that it has translated into legal codification, it is indeed a natural place to start for AI. In this article, we anticipate directions in what we believe will become a dominant and impactful forthcoming debate, namely, how to conceptualise the relationship between data protection and AI impact. We begin by discussing the value canvas i.e. the ethical principles that underpin data and AI ethics, and discuss how these are instantiated in the context of value trade-offs when the ethics are applied. Following this, we map three kinds of relationships that can be envisioned between data and AI ethics, and then close with a discussion of asymmetry in value trade-offs when privacy and fairness are concerned.

Keywords Artificial intelligence · Machine learning · Data · Data protection · Ethics · Audit · Impact assessment

1 Introduction

In the growing literature on artificial intelligence (AI) impact assessments, which includes technological auditing of metrics such as privacy, fairness and performance, and human rights, social and environmental impact assessments [1], the literature on data protection impact assessments (DPIA) is heavily referenced and drawn upon [2–5]. Given the relative maturity of the data protection debate and that it has translated into legal codification (most explicitly in the general data protection regulation (GDPR)) [6–9], drawing upon it is a natural place to start for AI. Indeed, when legality is referenced, the GDPR legislation is often mapped on to discussions regarding compliance of AI systems [3–5]. A paradigmatic example of this can be found in the UK's Information Commissioner's Office 'Guidance on AI and data protection' [5].

In this article, we anticipate directions in what we believe will become a dominant and impactful debate, namely how to conceptualise the relationship between data protection (which we read mainly as an expression of the value of

privacy) and AI impact (which we read predominantly as an expression of the value of fairness). We begin by discussing the **value canvas** i.e. *the principles that underpin data and AI ethics*, and discuss how these are instantiated in the context of value trade-offs when the ethics are applied. Following this, we map a triad of potential relationships between data and AI ethics:

- AI and data as pyramidal, with data protection being the foundation;
- consideration of AI impact will cause a renegotiation of data protection concerns and the two will be integrated in this process; and
- AI impact assessments should be independent of data protection because AI systems have unique challenges that are irreducible to data protection concerns.

We then close with a discussion of asymmetry in value trade-offs when privacy and fairness are concerned. Our key contributions to the debate are:

- **privacy as non-foundational** notwithstanding extensive legal frameworks of data protection as a fundamental right, at present there is no *philosophical* consensus that privacy is primary and a fundamental value, indeed, it may be argued that the value of privacy has no founda-

✉ Emre Kazim
e.kazim@ucl.ac.uk

¹ University College London, London, UK

tional status but rather is a derivative of modern political structures;

- **inequivalence of privacy in data and AI** where data governance i.e. data stewardship, deals with privacy and fairness in data, contrastingly, AI systems may introduce privacy concerns by their process i.e. privacy concerns may be raised however not the same kind of privacy issues as that which is provisioned for by data protection;
- **privacy vs. fairness** here the two values of privacy and fairness can be thought of as in direct conflict with one another. If the argument is sound with respect to the fundamental value of data protection being privacy and AI being fairness, the two are incompatible. If this is rejected and the assumption is made that the values must remain dynamic, then data and AI assessment will each need its own impact assessment.

2 Value canvass

In recent years, with the increased development and deployment of autonomous systems—popularly referred to as artificial intelligence (AI)—a growing concern has arisen as a result of high-profile cases of harm. The awareness of the social impact and ethical implications of AI has increased within the various stakeholders—namely, the academy, government, civil society (through NGOs) and industry. Examples of harm that were observed are bias in systems such as recruitment [10] and criminal justice sentencing (where particular demographics are prejudiced against), voter manipulation and misdiagnosis of cancer patients [11–13]. With these a growing consciousness developed within wider society and developers of these technologies that something needs to be done. Indeed, what is now referred to as ‘AI ethics’ or ‘trustworthy AI’ or ‘responsible AI’ is the body of literature that has resulted because of this consciousness and debate [14]. In our reading, the field of AI ethics has undergone three broad phases. The first was an AI ethics set of principles [15]. The second phase was an ethical-by-design approach, which was an engineering focused problem-solving exercise [16]. We read the third—indeed the current phase—as concerned with the need to standardise and operationalise the AI ethics discipline. Where we read assurance as a broad term to encompass certification and audit. The quest to achieve trustworthy AI has matured to the point where appropriate governance, regulation, impact assessment and auditability standards are being proposed and formulated [2–5, 17]. More broadly, we can define AI ethics in terms of applied ethics ‘the psychological, social and political impact of AI’ and in terms of human-centric AI ethics ‘the development and deployment of AI systems that respect human dignity and autonomy’ [18], p. 2, 5.

Table 1 Practical instantiation of the value canvass

Value	Practical instantiation
Respect for human dignity (fundamental human rights)	Human agency and oversight
Prevention of harm	Technical robustness and safety; Privacy and data governance; Societal and environmental well-being
Principle of explicability	Transparency
Principle of fairness	Diversity, non-discrimination and fairness; Societal and environmental well-being; Accountability

In contrast to the relatively ‘new’ AI ethics, with its associated impact assessment, debate and literature, is the more mature data protection impact assessment and data privacy literature and law. As discussed above, this maturity is most clearly demonstrated in GDPR legislation. Given this context, in this section, we explore the interrelation between concepts and terms with respect to data and AI ethics.

2.1 AI and data protection impact assessments

Both purport to be responses to moral concerns [14]. In the political and public debate, *the premise is that real and potential ‘harms’ can be mitigated through impact assessments* by ensuring that systems that use data and AI are evaluated in design and deployment with respect to a set of standards (standards that have been codified into law in the case of GDPR) [7, 19]. These standards are read as a translation or practical instantiation of values (or principles), which both drive and underpin the notions of ‘harm’ being mitigated.

These values/principles to practical instantiation—i.e. the translation of ethics into engineering practice and legal recourse—requires in the first instance an articulation of the values/principles (abstract philosophical ethics) into by-design and legal norms (applied ethics) [20]. An example of this can be found in the *European Commission’s Guidance on Trustworthy AI, which maps values for a system to be considered ‘Trustworthy’* (summarized below, in Table 1). Note that some practical instantiations map to multiple values [2].

2.2 Multiplicity of values

Elsewhere others offer analysis of the practical instantiation of values [18]; for the present, it suffices to state that *there is a value canvass that is concretised in engineering systems, legislation etc. i.e. from abstract to applied ethics*. The manner in which the value canvass is presented above, which is




Fig. 1 Context dependency of values in data and AI ethics

drawn from the current AI ethics/systems/standards debate, presents a multiplicity of values. This reflects the numerous values and aspects of an ethical (in the case of AI, ‘trustworthy’) environment of data and AI utilisation [21]. Figure 1 shows how values are shared and principles are expressed in data and AI ethics.

Nonetheless, within the literature, the nature of these values i.e. in how they interrelate, is seldom discussed. Instead, in the context of AI ethics, there is a *significant strand that accepts that the practical instantiation of ethical values and principles will require trade-offs* i.e. the principle of explicability (practically instantiated as AI explainability) and the principle of prevention of harm (which in one dimension is practically instantiated as AI robustness) may have to be negotiated [1–3, 14]. Typically, the response to this is a call for context sensitive trade-offs. For example:

- **privacy in the context of criminal justice sentencing may be justifiably traded for explicability**, which guarantees a person transparency with respect to how a sentence is calculated;
- conversely, explicability may be traded in systems that process social media data, where privacy preserving protocols may render it impossible to explicate with respect to processing specifically to the person [22–24].

2.3 Context sensitivity

Such trade-off debates are premised upon the idea that values are dynamic i.e. context dependent and are determined in context. Another way of stating this is that *values change in importance with respect to the context*. This premise is contentious from a philosophical perspective insofar as some values are considered ‘fundamental’—this is particularly true in terms of the legal codification of respect for human dignity i.e. human rights law. Here, the value of human dignity is explicitly stated in

terms of being non-negotiable (inviolable, inalienable, etc.). Another value that *prima facie* is considered non-negotiable is that of fairness i.e. non-discrimination (most acutely stated in terms of racial and gender equal treatment). Indeed, *Recital 1 of GDPR ‘Data Protection as a Fundamental Right’* states that:

‘The protection of natural persons in relation to the processing of personal data is a fundamental right’,

suggesting that similar to human dignity and fairness, privacy is a fundamental value (read: non-negotiable, inviolable, inalienable, etc.) [25]. The *Recital appears to equivocate respect for human dignity with respect for privacy*. Metaphysically, this is a ‘thick’ concept entailing significant argument: **at present there is no philosophical consensus that privacy is primary and a fundamental value**, indeed, it may be argued that the value of privacy has no foundational status but rather is a derivative of modern political structures [26, 27]. Expanding upon this, it is important to distinguish between discussions of privacy as a fundamental right *qua* law, and privacy as a fundamental right *qua* the philosophical tradition. Regarding the former, there is ample legal literature and jurisprudence that addresses the fundamental right of privacy [see, Universal Declaration of Human Rights (Article 12), the European Convention of Human Rights (Article 8) and the European Charter of Fundamental Rights (Article 7)]; however, and contrastingly, regarding the latter there is considerable contention [27–30]. Indeed, it is an open question as to whether or not a first-principles philosophical defence of privacy is required for privacy to be considered a fundamental (legal) right, or, if it is sufficient, as it most often the case, that the ‘fundamental right of privacy’ is derived from the notion of ‘human dignity’ [31, 32], see also GDPR (Article 88).

Drawing this together, we believe that *there is a lacuna in the literature concerning the expression of values that premise both data and AI debates*, and the notion of context dependent trade-offs begotten by the need to apply ethics practically (in engineering and legal terms). We believe that this lacuna, which can also be stated in terms of a contradiction, will be exposed further with the maturation of AI ethics and its practical instantiation in the form of impact assessments and how this relates to the more advanced discourse/standards/legislation on data.

3 Relating data and AI ethics and impact assessments

In this section, drawing upon the previous discussion, we map three broad ways in which data protection and AI impact assessments may relate.

3.1 AI impact assessment sits on top of data protection assessments

We noted in the introduction that AI ethics and AI impact assessment concerns draw upon the data ethics literature and data protection legislation. Indeed, this can be shown in a number of ways; the most clear being the reference to ‘data’ ethics and principles in AI ethics and impact assessment literature. This is most exemplified by the change in title from the UK ICO’s call for consultation on their ‘draft AI auditing framework guidance for organisations’ [33] to ‘Guidance on AI and Data protection’ [5], where they detail the building of AI concerns upon existing data ethics and legislation [34]. An alternative way to show this is by drawing upon UK and EU GDPR legislation, as well as UK-ICO and EU-Article 29 Data Protection Working Party explanatory documents. Here the following argument can be made: A DPIA is the process that is used to identify and thereby reduce the risks associated with the analysis of personal data. An appropriate DPIA will fulfil legal obligations as required by the GDPR and other data protection legislation. A DPIA is required when processing operations are ‘likely to result in high risk’ with respect to personal data. According to the European Data Protection Board (EDPB) on DPIAs (WP248rev01 Section 3 paragraph 8);

“Innovative use or applying new technological or organisational solutions, like combining use of finger print and face recognition for improved physical access control, etc. The GDPR makes it clear (Article 35(1) and recitals 89 and 91) that the use of a new technology, defined in “accordance with the achieved state of technological knowledge” (recital 91), can trigger the need to carry out a DPIA. This is because the use of such technology can involve novel forms of data collection and usage, possibly with a high risk to individuals’ rights and freedoms. Indeed, the personal and social consequences of the deployment of a new technology may be unknown.”

Note the example of ‘face recognition’ technology i.e. an AI system. More directly, the UK ICO guidelines on DPIA include explicit reference to AI technologies as innovative [35]. As such, an algorithmic system is classed as ‘likely to result in high risk’. An implication of this is that DPIA’s may be required when using any AI system.

Ergo, no new or fundamental differences are introduced by special consideration for AI systems. Data protection is seen as foundational and even as the overarching framework within which AI falls under. This is akin to reading broader ‘digital ethics’ as effectively data ethics and more specifically taking privacy to be the fundamental organising ethical principle and value of all new digital technology impact assessments.

3.2 Data protection assessments should be adapted and modified according to the specifics, and unique challenges, of AI systems

Here, data protection assessments are read as guiding the discussion and providing direction for the forthcoming AI impact assessments. This has three dimensions:

- **data protection impact assessments have been well developed and deployed** for some years now and as such there is considerable practical knowledge and literature in this field thereby making it an invaluable resource from which to draw upon;
- **data protection has been codified into laws** and as such future AI impact assessments should be developed in such a way as to cohere with these laws, hence it is pragmatically better to adjust data protection provisions rather than introduce a wholesale new framework;
- **AI systems rely on data processing** and as such the two are necessarily related and therefore their impact assessments are necessarily intertwined.

3.3 AI impact assessments should be independent of data protection because AI systems have unique challenges that are irreducible to data protection concerns and its fundamental value of privacy

Here, notwithstanding the necessary relationship between AI systems and data processing, a clear difference is envisioned between AI impacts and data assessments. The *motivation behind this relates to the claim that there are fundamental ethical implications and questions in AI that cannot be read as transferable from data protection*. For example, there may be a data set that is secured and pre-processed etc. in such a way as to satisfy data protection (and fairness in the dataset itself, etc.) and yet despite satisfying ethical norms and legal compliance on the ‘data’ front the AI system may nonetheless have ethical problems unique to how it operates.

Furthermore, data governance (also referred to as data stewardship), namely ‘defining, implementing and monitoring strategies, policies and shared decision-making over the management and use of data assets’ [36], p. 6. in the context of data protection is not equivalent to the discussions or responsibilities associated with human oversight mechanisms, with respect to decisions that an AI system may make. Indeed, with respect to AI governance and ‘stewardship’ human intervention is read in terms of keeping-the-human-in-the-loop, where decisions are reviewed and checked by humans (ensuring that responsibility falls clearly on humans i.e. non-solely automated decisions c.f. [3, 37]). Below we detail further discrepancies:

- Where data governance and stewardship deals with privacy and fairness in data, contrastingly, AI systems may introduce privacy concerns by their process i.e. privacy concerns may be raised however not the same kind of privacy issues as that which is provisioned for by data protection.
- **AI also introduces issues of opacity** i.e. the so-called black box, and the need for explainability and fairness: importantly it is a different kind of fairness problem to the one that can be addressed and satisfactorily solved through, for example, pre-processing of data. AI requires value judgments and this is a far more fluid and more subjective an intervention than data protection, which has a more stable form of fairness i.e. making the data set balanced. Therefore human oversight is more likely to be subjective and plural in AI systems comparatively than with respect to data stewardship judgements (where more straightforward compliance with privacy provisions and fairness is the case).
- **Finally—perhaps the key difference—an argument can be made that the fundamental value driving AI ethics (and its impact assessment) is fairness.** Stated negatively; the most ubiquitous risk of AI systems can be construed as bias, which can be read as in conflict with the fundamental value of privacy. Another way of stating this is that the guiding ethical principle is fairness—hence why explainability and accuracy is demanded from AI impact assessments—all of which may be at the cost of privacy [38].¹
- *Within Fairness* Here we take ‘fairness’ to be an umbrella under which debates on the nature of fairness is understood. For example, there is a vibrant debate on whether fairness entails equality of outcomes, or if equality of opportunity is fairness. Furthermore, what metrics should we identify as crucial to equality (gender, socio-economic class, demographic, etc.) [39]. Fairness as a value is discussed here in terms of justice and one notion of fairness is traded against another i.e. it depends on the notion of the collective good one ascribes to. This can also be stated as ‘political justice’.

These two—i. and ii.—can be read as trade-offs internal to each value; **ii. is different to i. because ii. is ‘closed’, in other words there is a plurality in fairness notions and these will have to be selected against (traded) with respect to one another.** For example, where gender parity may come at the cost of racial demographic fairness there will be a debate. Contrastingly, privacy may be thought of in terms of an absolute (full privacy protection is an idealised state). If we could have full privacy protection, then this would be the ideal state; however, we have practical reasons why this is not possible (full anonymisation may lead to amplification of bias or retard the right to withdraw consent, etc.); whereas, and contrastingly, with respect to fairness, the value choices are fundamentally incommensurable.

The type of moral deliberation with respect to privacy being traded for other values and the internal/plural intrinsic nature to the debates within the notion of fairness, are of a different class. Even if in both cases, the moral deliberation is presented in terms of trade-offs, this fundamental asymmetry exists.

Privacy vs fairness Here the two values of privacy and fairness can be thought of as in direct conflict with one another. **If the argument is sound with respect to the fundamental value of data protection being privacy and AI being fairness, the two are incompatible** and either a clear position would have to be taken on what is more important i.e. a hierarchy, which would allow integration of data protection and AI impact assessments through a framework that priorities one over the other. In other words, in cases where a value must be traded, always choose one over the other. If this is rejected and the assumption is made that the values must remain dynamic i.e. value trading as context specific and not settled theoretically/philosophically beforehand, then data and AI assessment will each need its own impact assessment and how the two relate will have to be worked out through another mechanism.

Finally, if no fundamental value judgement with respect to the priority of privacy (data) and fairness (AI), is made, then the two can be integrated; however, this will not be by mapping AI on to data protection provisions but by a negotiated approach that sees these two impact assessments

4 Asymmetry in privacy and fairness value trade-offs

In light of the possibilities presented and the discussion of values in the previous section, to flesh out the tension between data and AI ethics/standards in the context of impact assessments, in this section, we introduce further nuance on how value explication may operate in different contexts.

- *Privacy Traded* Here we refer to the concretisation of privacy in terms of a trade-off with things like accuracy, and the quality of a service. Privacy is traded against another value

¹ One way that may have solved both the problem of privacy and fairness in AI systems is through full anonymisation, however this was shown to simply carry through and indeed mask bias in data sets used by AI systems, thereby not addressing the problem and perhaps making it worse. As a result of this data minimisation has become the imperative, rather than anonymisation.

(grounded on different values) as fundamentally in tension. One consequence of this will be a reworking of data ethics and DPIAs in light of more and more sophisticated applications of AI systems.

5 Conclusion

In this article, we have presented a series of reasons as to why a move from data to AI ethics is unlikely to be straightforward and that attempts to map AI ethics to data protection provisions and data related legal codes will raise contradictions and sophisticated value judgements. This problem will be particularly acute when AI ethics built upon data ethics is applied practically—as is the case in the impact assessment literature. Indeed, the arguments we present should be read as *prima facie* reasons as to why considerable development of this interrelation (both theoretical and in terms of practical instantiation i.e. engineering practice and law), is needed. We hope this article stimulates this debate and is a step towards filling this lacuna.

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

1. Ada Lovelace and DataKind UK.: Examining the black box: Tools for assessing algorithmic systems. Technical report, AdaLovelace Institute, <https://ico.org.uk/media/about-theico/consultations/2617219/guidance-on-the-ai-auditing-framework-draft-for-consultation.pdf> (2020)
2. AI HLEG.: Ethics guidelines for trustworthy AI. B-1049 Brussels. (2019)
3. UK-ICO.: Guidance on the ai auditing framework: Draft guidance for consultation. Technical report, Information Commissioner's Office, United Kingdom. (2020)
4. AI-HLEG.: The Assessment List For Trustworthy Artificial Intelligence (ALTAI). (2020)
5. UK-ICO.: Guidance on AI and data protection. (2020)
6. <https://gdpr-info.eu/>
7. UK Data Protection Act 2018
8. UK-ICO.: Guide to data protection regulation. <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/> (2020) Accessed: 30 August 2020
9. Voigt, P., Axel, Von dem B.: The EU general data protection regulation (GDPR). A Practical Guide, 1st Ed., Cham: Springer International Publishing. (2017)
10. Schulte, J.: AI-assisted recruitment is biased. Here's how to make it more fair. World Economic Forum. <https://www.weforum.org/agenda/2019/05/ai-assisted-recruitment-is-biased-heres-how-to-beat-it/> (2019). Accessed 30 Aug 2020
11. Hao, K.: AI is sending people to jail—and getting it wrong. MIT Technology Review <https://www.technologyreview.com/2019/01/21/137783/algorithms-criminal-justice-ai/> (2019) Accessed 30 Aug 2020
12. Burkell, J., Priscilla, M., Regan, P.M.: Voter preferences, voter manipulation, voter analytics: policy options for less surveillance and more autonomy. *Internet Polic. Rev.* **8**(4), 1–24 (2019)
13. Challen, R., et al.: Artificial intelligence, bias and clinical safety. *BMJ Qual. Saf.* **28**(3), 231–237 (2019)
14. Jobin, A., Ienca, M., Vayena, E.: The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* **1**(9), 389–399 (2019)
15. Piano, S.L.: Ethical principles in machine learning and artificial intelligence: cases from the field and possible ways forward. *Human Soc. Sci. Commun.* **7**(1), 1–7 (2020)
16. Leslie, D.: Understanding artificial intelligence ethics and safety: a guide for the responsible design and implementation of AI systems in the public sector. The Alan Turing Institute. (2019)
17. Kazim, E., Koshiyama, A.: AI assurance processes . Available at SSRN: <https://ssrn.com/abstract=3609292> (2020)
18. Kazim, E., Koshiyama, A.: A high-level overview of AI ethics . Available at SSRN: <https://ssrn.com/abstract=3609292> or (2020). <https://doi.org/10.2139/ssrn.3609292>
19. European-Commission: White paper on artificial intelligence—a european approach to excellence and trust (2020)
20. Whittlestone, J., Nyrup, R., Alexandrova, A., Dihal, K., Cave, S.: Ethical and societal implications of algorithms, data, and artificial intelligence: a roadmap for research. Nuffield Foundation, London (2019)
21. Lukowicz, P.: The challenge of human centric AI. *Digitale Welt* **3**, 9–10 (2019)
22. Chouldechova, A.: Fair prediction with disparate impact: a study of bias in recidivism prediction instruments, arxiv pre-print. (2016)
23. Corbett-Davies, S., Goel, S.: The measure and mismeasure of fairness: A critical review of fair machine learning. arXiv pre-print (2018)
24. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. arXiv preprint arXiv: 1908.09635 (2019)
25. Palmirani, M., Martoni, M., Rossi, A., Bartolini, C., Robaldo, L.: Legal ontology for modelling gdpr concepts and norms. *JURIX*, 91–100 (2018)
26. Richardson, J.: Law and the Philosophy of Privacy. Routledge. (2015)
27. Negley, G.: Philosophical views on the value of privacy. *Law Contemp. Probs.* **31**, 319 (1966)
28. Moore, A.D.: Defining privacy. *J. Soc. Philos.* **39**(3), 411–428 (2008)
29. Thomson, J. J.: The right to privacy. *Philos. Public Aff.* 295–314 (1975)
30. The English Law of Privacy: An Evolving Human Right—Lord Walker. Supreme Court Justice Lord Walker of Gestingthorpe. (2010). https://www.supremecourt.uk/docs/speech_100825.pdf
31. Floridi, L.: On human dignity as a foundation for the right to privacy. *Philos. Technol.* **29**(4), 307–312 (2016)
32. Schoeman, F. D.: Privacy and social freedom. Cambridge university press (1992)

33. UK-ICO: ICO & Stakeholder consultation. <https://ico.org.uk/about-the-ico/ico-and-stakeholder-consultations/ico-consultation-on-the-draft-ai-auditing-framework-guidance-for-organisations/to> (2020) Accessed 30 Aug 2020
34. UK-ICO: Consultation responses. <https://ico.org.uk/media/about-the-ico/consultation-responses/2618057/ai-guidance-consultation-responses-20200730.pdf> (2020) Accessed 21 Oct 2020
35. UK-ICO: Guide to data protection: examples of processing likely to result in high-risk. <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/data-protection-impact-assessments-dpias/examples-of-processing-likely-to-result-in-high-risk/>. (2020) Accessed 21 Oct 2020
36. Data governance and data policies at the European Commission: European Commission, Secretariat-General. Available: https://ec.europa.eu/info/sites/info/files/summary-data-governance-data-policies_en.pdf (2020)
37. Wilkinson, M.D., et al.: The fair guiding principles for scientific data management and stewardship. *Sci. Data* **3**(1), 1–9 (2016)
38. UK-ICO: Data minimisation and privacy-preserving techniques in ai systems. (2020) Accessed 7 June 2020
39. Oneto, L., Chiappa, S.: Fairness in machine learning. In recent trends in learning from data, 155–196. Springer (2020)

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