AI and Ethics
https://doi.org/10.1007/s43681-022-00179-z

COMMENTARY

Proposed EU AI Act—Presidency compromise text: select overview and comment on the changes to the proposed regulation

Emre Kazim[1,2] · Osman Güçlütürk[3] · Denise Almeida[4] · Charles Kerrigan[5] · Elizabeth Lomas[4] · Adriano Koshiyama[1,2] · Airlie Hilliard[2,6] · Markus Trengove[2,7]

Received: 8 March 2022 / Accepted: 21 May 2022
© The Author(s) 2022

Abstract
With its proposed EU AI Act, the EU is aspiring to lead the world in admiral AI regulation (April 2021). In this brief, we summarise and comment on the 'Presidency compromise text', which is a revised version of the proposed act reflecting the consultation and deliberation by member states and actors (November 2021). The compromise text echoes the sentiment of the original text, much of which remains largely unchanged. However, there are important shifts and some significant changes. Our main comments focus on exemptions to the act with respect to national security; changes that seek to further protect research, development and innovation; and the attempt to clarify the draft legislation's stance on algorithmic manipulation. Our target readership for this paper is those who are interested in tracking the evolution of the proposed EU AI act, such as policy-makers and those in the legal profession.

Keywords Regulation · Compliance · Standards · Legislation · Artificial intelligence · Accountability · Governance · Fairness · Transparency

Eyes of a peacock's feather
Butterflies beating -
Winged blue and red
Spotted red and circled blue
Deep blue and blood red
Crimson and azure
Blurring into forest green
Just as an eye blinks and sees
Something more than what is there
- Tahsin

1 Introduction

With its proposed EU AI act (April 2021) [1], the European Union (EU) is aspiring to lead the world in artificial intelligence (AI) regulation. This proposed legislation has garnered much attention and is set to become the first, explicit, regulation by a major jurisdiction specifically addressing AI. Elsewhere, we have provided a summary and commentary of the first draft version of the act and refer readers to this piece for an overview [2]. In this brief, we summarise and comment on the 'Presidency compromise text' (November 2021) [3], which is a revised version of the proposed act reflecting the consultation and deliberation by member states and actors. While the original text remains largely unchanged, there are important shifts and some significant changes with respect to the scope of the legislation and member state autonomy (in particular with regard to national security) in the updated version. Our target readership for this paper is those who are interested in tracking the evolution of the proposed EU AI act, such as policy-makers and those in the legal profession.

✉ Emre Kazim
e.kazim@ucl.ac.uk

1 Department of Computer Science, University College London, Gower St, London WC1E 6EA, UK

2 Holistic AI, 18 Soho Square, London W1D 3QL, UK

3 Faculty of Law, Boğaziçi University, 34342 Bebek/Istanbul, Turkey

4 Department of Information Studies, University College London, Gower St, London WC1E 6EA, UK

5 CMS Cameron McKenna Nabarro Olswang LLP, Cannon Place, 78 Cannon Street, London EC4N 6AF, UK

6 Institute of Management Studies, Goldsmiths, University of London, New Cross, London SE14 6NW, UK

7 School of Public Policy, University College London, 29 Tavistock Square, London WC1H 9QU, UK

Published online: 27 June 2022

We note that the major changes can be grouped into exemptions for national security, research and development, and general purpose systems; clarifications in regard to providing a clear definition of AI, critical infrastructure, personal data, non-personal data, and manipulation; extensions of prohibitions to include private sphere social scoring and socio-economic vulnerability; the explication of the use of AI in insurance as high-risk; clarification on the scope of the proposed act in terms of the removal of jurisdictional boundaries; and the classification of AI systems already mandated to undergo a third-party conformity assessment as high risk. In Sect. 2, we elaborate on each of these key changes and provide further context through a discussion of each point. In Sect. 3, we identify some key themes in the updated compromise text and provide a commentary of each theme, which is the major contribution of this paper. Our main comments focus on:

- *Security:* given that much of the societal concern regarding AI has traditionally been with respect to the use of AI in the security services (e.g., concerns about the movement towards a surveillance state or its use by law enforcement agencies), we believe that this exempting of AI in the context of National Security is a significant change and is likely to foundationally alter how the Text of the original AI Act had structured the relationship of trust and AI with respect to citizens and consumers.
- *Research, development and innovation:* although we called for this change in our initial response [2], we have two further points of contention to raise. The first concerns the possibility that this can be gamed in the form of structuring work as research and development, when in fact it is work that is developed with a particular purpose in mind (which can readily be masked). The second concerns expanding this exemption to small- and medium-sized enterprises (SMEs), where an analogue of this exemption could be made for an SME, defined in terms of some monetary threshold that ensures the majority of, for instance, start-ups are exempt. This is critical as the compromise text seems to imply that SMEs may play a significant role in driving tech research and development (R&D) and innovation, and this seems to evidence historic assumptions about how R&D and innovation plays out today.
- *Manipulation:* we recognise this attempt to clarify; however, the clarification is just as ambiguous as the original phrasing. In fact, whilst there is legal precedent around it, the insertion of the term "reasonably" could create an additional layer of ambiguity. The intuition and intent are clear—the legislation seeks to address the "nudging" and manipulation that is much of the focus of the AI ethics literature as it relates to potential psychological control. However, we believe that the inability to capture this in the language of the legislation is likely to be a result of the work to be done conceptually and empirically in this area.

Note that readers will require familiarity with the original version of the act. The numbers in brackets (..) refer to the page number of the Presidency compromise text [3].

2 Summary of changes

In this section, we summarise the main changes to the legislation. This is primarily a condensed version of the compromise text's own summary of the 'Main Changes' (3–5).

- *Exemptions:* the compromise text has been revised in such a way as to provide specific exemptions to the proposed legislation. These are:
    - *National security:* exemptions with respect to National Security are perhaps the most significant change to the legislation. The revision states that there should be: "an explicit reference to the exclusion of national security from the scope of the proposed regulation […] national security remains the sole responsibility of each Member State" (3). The problem with the term "national security" is that it is extremely vague. What could or should be considered as a matter of national security? The fact that making this determination is up to each Member State's discretion creates the risk of diverging applications, which could undermine the goal of having a harmonised regulation of AI within the EU. Despite the fact that it is quite common for European Regulations to include exemptions for National Security (e.g. the General Data Protection Regulation includes such provisions [4]), these are usually very well defined and limited to specific scenarios, such as the limitation of specific user rights. This is increasingly difficult to define and could lead to states relying on these exemptions to deploy mass surveillance solutions on their citizens. We see this as a major area to define: how does National Security operate and interact with the protections and rights of individuals? We also see it as notable that the EU takes a different approach to the US. The Final Report of the US National Security Commission on Artificial Intelligence takes an assertive position in relation to these technologies based on a view that: "We fear AI tools will be weapons of first resort in future conflicts". Cyber and other uncon-

ventional operations in the conflict in Ukraine are likely to mean that EU policy-makers reconsider their position on national security and AI.

- *Research and development:* the original proposal was ambiguous with respect to whether research and development would fall under the remit of the legislation. With this, there was likely to be confusion and burden on institutions, such as universities and small–medium-sized enterprises. As such, the explicit exemption—"it has been clarified […] that the AIA [Algorithm Impact Assessment] should not apply to AI systems and their outputs used for the sole purpose of research and development." (3)—is indeed an important clarification and we welcome this. The compromise text provides a thorough statement of this exemption (11): "It is therefore necessary to exclude from its scope AI systems specifically developed and put into service for the sole purpose of scientific research and development and to ensure that the regulation does not otherwise affect scientific research and development activity on AI systems." As a general point, commercialisation of R&D can often be supported by a "sandbox" approach. However, this may open up debate about the need for and role of a regulator (and, at this early stage in the policy discussion it would be helpful to explore the need for a regulator alongside the terms of regulation.
- *General purpose systems:* general purpose systems have been exempted; AIAs are only triggered when the systems are used in a particular context i.e. for a particular purpose (5; see also the introduction of Article 52a (68)). The inserted text is paragraph 70a (26), within which the following is stated: "[…] Therefore, the placing on the market, putting into service or use of a general purpose AI system, irrespective of whether it is licensed as open-source software or otherwise, should not, as such, trigger any of the requirements or obligations of this regulation". We read this as a corollary to the above-noted exemption for research and development. Similarly, we welcome this exemption as it is difficult to see how an impact assessment, which assesses risk (primarily with respect to harm upon humans, as is constituted), can be carried out without considering the use case. However, it is also worth considering that the purpose is to leave these "general purpose" systems outside the scope unless they are used for an "intended purpose", which is defined, however, once again, with broad terms. Interestingly, the text clearly does not want to leave these completely outside the scope, but details are uncertain, and we read this as confusing from the regulatory perspective. Under Article 3(12), these have been deemed not to have intended purpose within the meaning of AIA. What is the implication of this? This is not clear as "having an intended purpose" is not a prerequisite for something. High-risk systems in Annex III are drafted with reference to the intended use but prohibited uses are not. When does a system become a general purpose system? What other provisions shall apply to these? Article 52a stipulates that the placing on the market, putting into service, or use of these shall not make these subject to the provisions of AIA. General purpose systems are not defined in the act but superficially explained under Recital 70a, which includes generative systems. Let's take a Generative Adversarial Network (GAN) as an example (used to create synthetic data, deep fakes, etc.) [5]. Deep fake creators are not general purpose in our opinion, but they could be considered as such according to Recital 70a. The term is confusing and should be delineated meticulously. In addition, almost every controversial system is built on one of these. We understand why someone who is not the developer would not focus on the commercial user or end-user, but just the users of an off-the-shelf AI system. Still, the main risk is intertwined with these "general purpose AIs". Even if they are exempted from some major requirements, they should be monitored during their lifecycle. In fact, these may be the correct regulatory point for an efficient regulation given that there are many AI systems that are just derivatives of these in some form. This concept, its definition, and the legal implications of the current version merit their own article, and as such we will close our comment here.

- *Clarifications:* a number of clarifications are offered, mainly with respect to terms, to increase legal clarity.
  - *Definition of AI:* a widely commented upon aspect of the original proposal was that the definition of AI was effectively all-encompassing [6]. This is a view that was grounded on the loose categorisation of systems (including 'statistical systems'). The compromise text has tried to tighten the definition so that there is more legal clarity and to exclude software that traditionally would not normally be considered AI (which we read in terms of legislation that covered things such as cyber security, etc.). As such, this tightening of the definition is something that a revision (such as the *Presidency compromise text*) would be expected to pick up on. The compromise text reads (italics ours):

'*artificial intelligence system' (AI system) means a system that*

(i) receives machine and/or human-based data and inputs,

(ii) infers how to achieve a given set of human-defined objectives using learning, reasoning or modelling implemented with the techniques and approaches listed in Annex I, and

(iii) generates outputs in the form of content (generative AI systems), predictions, recommendations or decisions, which influence the environments it interacts with;

Note that the Annex I referred to remains unchanged. Despite the attempt to tighten the definition, according to our reading, the definition is still very much loose and can still be read as all-encompassing. Indeed, the compromise text removed the reference to "software", and we cannot see why as it was an accurate component from which risks associated with AI arise. From a legal perspective, a clear and effective definition should be focusing on features differentiating AI systems from traditional software, which are mostly regulated via established legal frameworks.

- *Definition of critical infrastructure*: the original text, which defined critical infrastructure as "safety components in the management and operation of road traffic and the supply of water, gas, heating and electricity", is expanded to include "digital infrastructure". We recognise this as reasonable given the ubiquitous nature of digital infrastructure such as the internet (18; see also 5, 37). Comparing recent legislation in the UK, the National Security and Investment Act allows the UK government to scrutinise and intervene in certain acquisitions in "sensitive areas of the economy" (including artificial intelligence). This approach looks at means rather than ends. It will live alongside, and overlap with, the EU proposals.
    - *Definition of personal data:* This is clarified by referring to GDPR: "'personal data' means data as defined in point (1) of Article 4 of Regulation (EU) 2016/679" (37)
    - *Definition of non-personal data*: this is clarified by referring to GDPR—'non-personal data' means data other than personal data as defined in point (1) of Article 4 of Regulation (EU) 2016/679. Once ratified, the currently draft Data Act [7] should also provide further scope limitations on this category of data.
    - *Manipulation:* an attempt is made to clarify the notion of algorithmic manipulation. The language is changed from "AI systems intended to distort" to "AI systems materially distorting" human behaviour. It then introduces the idea of "reasonably" likely to cause harm, including "harms that may be accumulated over time". It is likely that this was introduced because of the significant ambiguity in the original text.
- *Extending prohibitions:*
    - *Private sphere social scoring:* the banning of social scoring by public authorities is extended to the private sphere (4). Although this does appear to extend a ban to personal analysis within companies (for example, for the purpose of marketing to customers), the amendments then clarify that "this prohibition should not affect lawful evaluation practices of natural persons done for one or more specific purpose in compliance with the law" (12), which reintroduces the ambiguity. In fact, this clarification is confusing since it seems to be referring to cases where a social scoring system that is covered by the AI Act could be lawful. However, this does not seem possible considering that social scoring applications falling under the scope of Article 5 are the ones that are leading to some kind of detrimental or unfavourable treatment of certain natural persons or groups.
    - *Socio-economic vulnerability:* the banning of manipulation, or in other words the special protection afforded to vulnerable groups (such as the elderly, the disabled, etc.), has been expanded to include vulnerability based on social and economic reasons (38). This insertion is useful as it allows the AI Act as well as regulators to capture currently unforeseen cases of manipulation to expand the scope of protection. However, it also creates significant legal uncertainty given that "social and economic situation" is a concept that is open to broad interpretation.
- *Insurance as high risk:* AI systems used in insurance have been added to the list of high-risk systems. The Compromise Text (19) reads: "AI systems are also increasingly used in insurance for premium setting, underwriting and claims assessment which, if not duly designed, developed and used, can lead to serious consequences for people's life, including financial exclusion and discrimination". Similar to our more general point regarding the ambiguity of definitions, although we recognise that including insurance is appropriate, does the inclusion of another case study signal that the regulation would have to count every single application (in the high-risk context)? The risk-based approach and determination of high-risk AI systems may have to be further delimited. Indeed, in terms of high-risk systems, it is interesting that they have taken out the research component of AI and justice

systems. This can nevertheless influence how case law is read and interpreted. More generally, it is likely that in practice "high risk" is too broadly framed. We foresee that the test will need to be by reference to the task set for an AI rather than simply the existence of any AI in the loop.

- *Scope:* although the change in the scope is not included in the 'main changes', we read the shifts in the text as a major change. The compromise text is changed in the following way: "In light of their digital nature, certain AI systems should fall within the scope of this regulation even when they are neither placed on the market, nor put into service, nor used in the Union. This is the case for example of an operator established in the Union that contracts certain services to an operator established outside the Union in relation to an activity to be performed by an AI system that would qualify as high-risk (10). We note this as a major change as the EU is effectively removing jurisdictional questions, which is a move that de facto makes the legislation global [8]. This is further augmented by the introduction of specific instructions for "importers and distributors of AI systems" (32). Additionally, there is an introduction to Article 14, which concerns transparency obligations, that requires disclosure of where the system is intended to be used, "inclusive of the specific geographical, behavioural or functional setting within which the high-risk AI system is intended to be used" (46). Finally, there appears to be a big political point for the UK here. Can the UK be a leader for AI standards (as its government hopes) if EU standards will necessarily apply extraterritorially to the UK? In that case, either the UK must be a rule follower or must require UK firms to comply with domestic and EU standards where those firms have any connection with the EU markets.
- *Mandatory third-party conformity assessment:* in the context of the treatment of high-risk systems (Article 6), a significant introduction to the compromise text is two paragraphs that stipulate that if an AI system is already mandated to undergo a third-party conformity assessment, it shall be considered high risk.

"An AI system that is itself a product covered by the Union harmonisation legislation listed in Annex II shall be considered as high risk if it is required to undergo a third-party conformity assessment with a view to the placing on the market or putting into service of that product pursuant to the above-mentioned legislation" (41).

"An AI system intended to be used as a safety component of a product covered by the legislation referred to in paragraph 1 shall be considered as high risk if it is required to undergo a third-party conformity assessment with a view to the placing on the market or putting into service of that product pursuant to above-mentioned [sic] legislation [...]." (41).

For a discussion of how conformity assessments of high-risk systems may lead to a Europe-wide ecosystem that leads to widespread auditing of systems, which we have previously surveyed the need for [9], see Mökander et al [10].

3 Comments

- *Security:* considering that the majority of the concerns from general society in regard to AI have traditionally been with respect to the use of AI in the security services (surveillance state, etc.), we believe that this exempting of AI in the context of National Security is a significant change that is likely to foundationally alter the relationship between trust and AI for citizens and consumers. Although not unexpected, as it aligns with the EU's understanding of national security and current legislative exemptions in other regulations with similar reach, the earlier divergence showed a potential desire to move to harmonisation in this area. Prior to this exemption, the relationship of trust was unambiguously one where the citizen and consumer were primarily protected by ensuring that particularly 'high-risk' systems would be banned and others would be governed by robust risk-governance provisions (e.g. Reporting and documentation). With this change, the relationship of trust and protection is shifted from procedure and transparency to citizens to the state itself. It is ironic that the change moves the EU proposal from one of the most progressive, citizen-centric, interventions, to one which imposes a strong trust on the state with respect to algorithmic systems. Indeed, this is further expanded upon with discussion of the military use of AI, where new text is introduced to assert that this should fall within existing international law, which it is implied, is sufficient to safeguard from risk (11; see also, 32).
- *Research, development and innovation: research, development and innovation:* although we welcome this change and had called for it in our initial response to the original draft [2], we have two further points. The first is that there is the possibility that the legislation may be gamed by those who dishonestly structure their work as research and development to mask the particular purpose of the work. Indeed, the requirement that the scientific research and development should be the sole purpose as per Article 2(6) and the emphasis that the research and development activity surrounding AI systems shall not be affected, in so far as such activity

does not lead to or entail placing an AI system on the market or putting it into service pursuant to Article 2(7), can be read as a recognition of such risk. However, these provisions introduce their own uncertainties with respect to the enforcement. First, the determination of which AI systems are for the "sole purpose of scientific research and development" is not an easy task. Second, it is uncertain how it could be ensured that an AI system developed under this exemption is not eventually placed on the market or put into the service. This is not too critical a point and reads somewhat conspiratorial, rather we raise it merely as a possibility. The second point concerns expanding this exemption to SMEs. An analogue of this exemption could be made for SMEs, which are defined using a monetary threshold that ensures the majority of, for instance, start-ups are exempt. On this point, the changes to the proposed legislation are minor and, according to our reading, inconsequential (for example they replace the term "small scale" with "SME", with no substantive difference to the scope (27; see also 33, 34). Indeed, despite the introduction of *Title V: Measures in Support Of Innovatio*n (69), there is little change and much of our criticisms go unanswered [2]. On a more general point, this is critical as the text implies that there is a significant role for SMEs in driving tech R&D and innovation, which evidences historic assumptions about how R&D and innovation plays out today.

- *Manipulation:* although an attempt at clarification has been made, the phrasing in the compromise text is just as ambiguous as the phrasing in the original draft. In fact, the insertion of the term "reasonably", whilst there is a whole case law around this term, could create an additional layer of ambiguity. The intuition and intent are clear—the legislation seeks to address the use of AI in "nudging" and manipulation, which is a major focus of the AI ethics literature as it links to psychological control. However, the lack of explication of this in the language of the compromise text is likely to signal the need for work to be done conceptually and empirically in this area. Alongside this, thought should be given to how interested parties will determine "reasonableness" in practice. Ultimately, the word will be for judges or regulators to define, but they will need guidance in the absence of precedent and, meanwhile, developers and users will be in the dark. To render the EU AI Act effective and enforceable, we should focus on more practical questions: What exactly is being manipulated? When, and under what conditions, is a manipulation materially distorting a person's behaviour? What qualifies as harm? How is algorithmic manipulation different from traditional forms of marketing or propaganda?

4 Conclusion

In this brief, we have selectively summarised the changes to the proposed EU AI act and have offered comments according to areas we think are of most significance. In closing we offer more speculative points on intellectual property (IP), the relationship between data and AI, and security:

- *IP:* readers may argue that the proposed regulation is not a complete overview of all AI considerations—this is a point we would concede. In particular, this is true of IP considerations, which we read as to be determined elsewhere in the EU's legislative agenda.
- *Regulator*: the regulatory frameworks need to be determined too as there are issues that relate to whether data protection and AI will fall under the same regulator. It is a step forward to articulating aspects of AI development and delivery; however, the interrelationship between data and AI is undeveloped [11].
- *Security:* although the approach to National Security exemptions is more in line with existing EU regulations, such as the GDPR, there is still a concern around the potential risks to citizens and their freedoms by the use of unchecked AI, such as facial recognition technologies (FRT), by law enforcement [12]. This is not a surprising position by the EU, but one which requires caution and the definition of boundaries and checks and balances around these exemptions.

Author contributions EK, OG, DA, CK lead the writing and conceptualization of the paper. AK was involved in the conceptualization of the paper. EL commented and workshopped the subject matter. AH and MT workshopped the paper. AH has edited the text.

Declarations

Conflict of interest The authors have no interests to declare.

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. European commission: 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, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206, Accessed 07 Dec 2021
2. Kazim, E., Kerrigan, C., Koshiyama, A.: EU proposed AI legal framework. SSRN Electron. J. (2021). https://doi.org/10.2139/ssrn.3846898
3. Council of the European union: 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-Presidency compromise text. (2021)
4. Brown, N.: understanding the data protection Act's national security and defence exemption-decoded. legal: internet, telecoms and tech law decoded., https://decoded.legal/blog/2021/04/understanding-the-data-protection-acts-national-security-and-defence-exemption, Accessed 04 Mar 2022
5. Koshiyama, A., Firoozye, N., Treleaven, P.: Generative adversarial networks for financial trading strategies fine-tuning and combination. https://arxiv.org/abs/1901.01751 [cs, q-fin, stat]. (2019)
6. Bryson, J.J.: europe is in danger of using the wrong definition of AI, https://www.wired.com/story/artificial-intelligence-regulation-european-union/, (2022)
7. European commission: data act: proposal for a regulation on harmonised rules on fair access to and use of data | shaping Europe's digital future, https://digital-strategy.ec.europa.eu/en/library/data-act-proposal-regulation-harmonised-rules-fair-access-and-use-data, Accessed 04 Mar 2022
8. Cailean Osborne: the European commission's artificial intelligence act highlights the need for an effective AI assurance ecosystem-centre for data ethics and innovation blog, https://cdei.blog.gov.uk/2021/05/11/the-european-commissions-artificial-intelligence-act-highlights-the-need-for-an-effective-ai-assurance-ecosystem/, Accessed 17 May 2022
9. Koshiyama, A., Kazim, E., Treleaven, P., Rai, P., Szpruch, L., Pavey, G., Ahamat, G., Leutner, F., Goebel, R., Knight, A., Adams, J., Hitrova, C., Barnett, J., Nachev, P., Barber, D., Chamorro-Premuzic, T., Klemmer, K., Gregorovic, M., Khan, S., Lomas, E.: Towards algorithm auditing: a survey on managing legal, ethical and technological risks of AI, ML and associated algorithms. SSRN Electron. J. (2021). https://doi.org/10.2139/ssrn.3778998
10. Mökander, J., Axente, M., Casolari, F., Floridi, L.: Conformity assessments and post-market monitoring: a guide to the role of auditing in the proposed European AI regulation. Mind. Mach. (2021). https://doi.org/10.1007/s11023-021-09577-4
11. Kazim, E., Koshiyama, A.: The interrelation between data and AI ethics in the context of impact assessments. AI Ethics. **1**, 219–225 (2021). https://doi.org/10.1007/s43681-020-00029-w
12. Almeida, D., Shmarko, K., Lomas, E.: The ethics of facial recognition technologies, surveillance, and accountability in an age of artificial intelligence: a comparative analysis of US, EU, and UK regulatory frameworks. AI Ethics. (2021). https://doi.org/10.1007/s43681-021-00077-w

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