

Published in final edited form as:

Nat Mach Intell. 2021 September ; 3(9): 738–739. doi:10.1038/s42256-021-00386-z.

Regulating AI in medicine in the United States and Europe

Kerstin N. Vokinger^{1,2}, Urs Gasser^{2,3}

¹Institute of Law, University of Zurich, Zurich, Switzerland

²Berkman Klein Center for Internet and Society, Harvard University, Cambridge, MA, USA

³TUM School of Social Sciences and Technology, Technical University of Munich, Munich, Germany

Abstract

The regulatory landscape for artificial intelligence (AI) is shaping up on both sides of the Atlantic, urgently awaited by the scientific and industrial community. Commonalities and differences start to crystallize in the approaches to AI in medicine.

In the United States, the US Food and Drug Administration (FDA) published a regulatory framework for AI applications in medicine in April 2019 and an action plan in January 2021 (ref. ¹). The FDA's leadership role in formulating regulatory guidance is a manifestation of the broader US national approach to the regulation of AI. In contrast to the European Union (EU), the US policy sustains from broad and comprehensive regulation of AI and instead delegates responsibilities to specific federal agencies, with an overarching mandate to avoid overregulation and promote innovation².

The EU approach to the regulation of AI applications in medicine (AIM) promises to be thicker by combining sector-specific with cross-sectional regulations. Most notably and with regard to the latter, the European Commission recently published a proposal for a robust legal framework for AI. The goal of the so-called AI Act is to promote the uptake of AI and the development of an ecosystem of trust³. The proposed legislation is not limited to medicine but presents a comprehensive risk-based regulatory approach to AI, identifying medicine as a high-impact sector for AI³. AIM need to fulfill the requirements of the AI Act as well as those already set forth under the existing EU Medical Device Regulation.

Despite the significant differences in the approach to AI regulation across the Atlantic, the United States and the EU share the goal to strengthen their position in the development and implementation of AI, with medicine as a key application area. Similarities—as well as differences—do not only exist at a fundamental level, but also become visible when zooming in on three focus areas of transatlantic regulation of AIM: (1) lifecycle regulation; (2) algorithmic bias; and (3) transparency to users.

Correspondence to: Kerstin N. Vokinger; Urs Gasser.

kerstinnoelle.vokinger@usz.ch; urs.gasser@tum.de.

Competing interests

The authors declare no competing interests.

Lifecycle regulatory framework

One of the greatest benefits of AIM resides in the ability to learn from real-world use, and the capability to improve their performance, a feature referred to as continual learning^{1,4}. Continual learning is a technique in which the decision logic of mathematical models is updated through new data while retaining previously learned knowledge. By contrast, locked AIM are trained on specific datasets. They often perform well on similar data, but could perform poorly in scenarios that are rare in the training process⁵.

Continual learning makes AIM unique among medical devices, also with regard to its regulation, which so far was not designed for adaptive systems. In response, both US and EU regulators propose a lifecycle regulatory framework that spans across all phases of AIM, from premarket development through postmarket performance, and takes into account the iterative nature of such learning systems^{1,3}.

Both the United States and EU introduce the possibility of predetermining changes to AIM and their performance at the moment of the initial authorization. The FDA proposes a so-called predetermined change control plan for the authorization of AIM. This plan shall include ‘what’ aspects the manufacturer intends to change through learning, as well as ‘how’ the algorithm will learn and change while remaining safe and effective¹. Changes may involve improvements in performance or changes in the indications for use—for example, expansion to a new patient population for which there had been insufficient evidence available to initially support that indication for use. However, there are many scenarios for which a new authorization is indicated. For instance, it would not be appropriate to allow AIM within the predetermined change control plan to develop from initially low-risk AIM to high-risk AIM (for example, leveraging the use of skin images to manage the healing of scars to the diagnosis of melanoma)⁶. Under the European regime, the establishment of a quality management system is also mandatory. Manufacturers are required to document, among other things, their strategy for managing modifications in their AIM, techniques for quality control or testing, and validation procedures³.

The FDA states, without further specification, that reporting of postmarket real-world performance shall maintain and assure safety and effectiveness of adaptive AI systems⁶. The EU regulator, in contrast, stipulates more detailed requirements for manufacturers and asks them to implement and maintain a postmarket monitoring system³. This system shall actively and systemically collect, document and analyse relevant data provided by users or collected through other sources on the performance of AIM throughout their lifetime, with the goal that the possible risks emerging from AIM can be addressed more efficiently³. Furthermore, AIM have to perform consistently throughout their lifecycle and meet an appropriate level of accuracy and robustness in accordance with the acknowledged state of the art³.

Compared to the United States, the AI Act in the EU stipulates more rigorous and more detailed prerequisites, and manufacturers carry greater responsibility. This approach will most likely increase the manufacturers’ workload and efforts, but may also enable trust in AIM.

Algorithmic bias

Bias is not an issue exclusive to AIM. Access to healthcare varies by factors such as gender, race and socio-economic status⁷. For example, clinicians may incorrectly discount the diagnosis of myocardial infarction of women because these patients are more likely to present with atypical symptoms⁸. As AIM are developed based on the collection and training of data from historical datasets, they may perpetuate biases present in the data and this may lead to wrong outcomes for certain population groups. Thus, an AI algorithm that learns from historical electronic health record data may not, for example, recommend testing for cardiac ischemia for women, delaying potentially lifesaving treatment⁹.

US and European regulators alike emphasize the importance of mitigating bias in AIM. The FDA highlights that such systems must be well suited for a racially and ethnically diverse intended patient population¹. Again, the European approach is more specific when compared to the United States. The proposed AI Act requires that the data used for AIM must meet quality criteria and be subject to appropriate data governance. For example, the data shall be examined in view of possible biases, and the data shall be relevant, representative, free of errors and complete, also with regard to the patients to which the AIM are intended to be applied. Furthermore, such data shall take into account, if indicated, the characteristics that are particular to the specific geographical, behavioural or functional setting within which the AIM is intended to be used³. While it is true that algorithms trained primarily on patient data from certain geographies may generalize poorly when implemented in new geographies due to different economic, socio-economic or ethnic features¹⁰, the EU regulatory approach may lead to implementation challenges. First, the EU is a region composed of multiple member states with ethnic representation differing within and among them. When a medical device is authorized in one EU member state, it may be marketed also in another without an additional authorization¹¹. Thus, a balanced geographical representation in one EU country may not necessarily be representative for another. Second, a requirement that domestic and international manufacturers of AIM have to (re-)train their algorithms on cohorts from European datasets or datasets that represent the European society may drive up costs and prolong the authorization process, with negative ramifications for innovation.

Transparency to users

So far, the US regulatory system has set the standards in terms of transparency regarding the authorization and application of a medical device. In particular, the FDA publishes summaries or statements for each approved medical device¹¹. By contrast, such data are not publicly available in Europe. The European Commission's database on medical devices (Eudamed2) is not publicly accessible.

The FDA recognizes that transparency is especially important for AIM and states that users should be informed about issues including usability, equity, trust and accountability of the AIM to ensure that users understand the benefits, risks and limitations of these systems. Currently, the FDA plans to hold a public workshop to elicit inputs from the broader community on how device labelling supports transparency to users¹.

The proposed EU AI Act is more specific. AIM shall be designed and developed in such a way to ensure that their operation is sufficiently transparent to enable users to interpret the system's output and use it appropriately. They shall be accompanied by instructions for use and information that include the characteristics, capabilities and limitations of performance, including the intended purpose, level of accuracy, robustness and cybersecurity, specifications for the training, validation and testing data, and the expected lifetime³. Furthermore, a public EU-wide database for AIM will be established. Manufacturers shall be obliged to provide meaningful information about their AIM. This registration will enable authorities, users and other interested people to exercise enhanced oversight over AIM and to promote trust³.

From a public interest perspective, the transatlantic emphasis on enhanced transparency is laudable. The proposed AI Act, if enacted as law, might point towards a new global gold standard in this respect. A public database, as proposed in the draft legislation, is a crucial step to improve trustworthiness by enabling a more robust evaluation of the benefits and risks of AIM. The transparency requirements set forth in the AI Act might also inform the evolution of US practices, which can benefit from an upgrade given that publicly available summaries of AIM authorization often lack a level of clarity and comprehensiveness that seems desirable.

Sharing core values

The United States and the EU address the same distinctive challenges of AIM through lifecycle regulation of (adaptive) AIM, highlight the importance of mitigating algorithmic bias and promote transparency. Despite these conceptual similarities, a comparison reveals important differences. Overall, Europe takes a more heavy-handed approach to the regulation of AIM, while the US approach emphasizes innovation and is more principle- and less detail-oriented. These differences are likely to lead to different outcomes in terms of innovation and adoption rates of AIM, whether in the medical context or generally.

Perhaps even more importantly, however, the regulatory differences may amplify previous policy decisions. For example, many algorithms at the core of AIM are trained on electronic health data¹². In the US, electronic health records have been introduced earlier compared to most EU countries, and more data such as demographic information, diagnoses, medications, medical procedures or survey results from self-reported questionnaires have been collected¹³. In Europe, by contrast, electronically collected health data are often less comprehensive¹⁴. Combined with the proposed regulatory requirements, it may be more challenging for manufacturers to obtain sufficient health data in Europe to train AIM in order to obtain authorization in the EU.

Despite the different dynamics and trajectories, it is important to remember that the United States and the EU share a set of core values and principles. For this reason, more information exchange and collaboration among the US and Europe authorities, as well as between research communities, seems desirable to strengthen the successful development of AIM for the sake of patient care and society at large—a goal that unites the United States and Europe.

Acknowledgements

This study was partially funded by the Swiss National Science Foundation (SNSF).

References

1. Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan. FDA; 2021.
2. Executive Office of the President: Guidance for Regulation of Artificial Intelligence Applications. The White House; 2020.
3. Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). European Commission; 2021.
4. Vokinger KN, Feuerriegel S, Kesselheim AS. *Lancet Digit Health*. 2021; 3 :337–338.
5. Oren O, Gersh BJ, Bhatt DL. *Lancet Digit Health*. 2020; 2 :486–488.
6. Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD). FDA; 2019.
7. Lavizzo-Mourey RJ, Besser RE, Williams DR. *NEJM*. 2021; 18 :1681–1684.
8. Canto JG, Rogers W, Goldberg RJ. *JAMA*. 2012; 8 :813–822.
9. Parikh RB, Teeple S, Navathe AS. *JAMA*. 2019; 322 2377 [PubMed: 31755905]
10. Kausha A, Altman R, Langlotz C. *JAMA*. 2020; 324 1212 [PubMed: 32960230]
11. Muehlematter UJ, Daniore P, Vokinger KN. *Lancet Digit Health*. 2021; 3 :195–203.
12. Rajkomar A, et al. *NPJ Digital Med*. 2018; 1 :18.
13. Toscano F, O'Donnell E, Unruh M, et al. *Eur J Publ Health*. 2018; 28 cky213.401
14. [accessed 13June 2021] Die elektronische Patientenakte. (Bundesministerium für Gesundheit <https://www.bundesgesundheitsministerium.de/elektronische-patientenakte.html>)