

The Regulation of Artificial Intelligence in the EU

Backgrounder

Whether and how to regulate AI has become a lively discussion all over the world. Only recently, some rather specific regulations have come into force in a number of countries including France, Germany, China and Canada. Elsewhere, a wave of legislative proposals are now under consideration. They are being intensely debated in Brazil, the United States of America, and the European Union. This backgrounder provides information about what is at stake when regulating artificial intelligence.




The Regulation of Artificial Intelligence in the EU [licence infos](#)

What does regulation of artificial intelligence mean?

Technologies that are important for society due to the risks and opportunities they present, are often subject to regulation, i.e. a set of rules governing a certain situation. The adoption of new technologies in society can have severe consequences that call for implementing new rules. Given that this involves settling disputes and conflicts in society, these rules are often enshrined in the law. Starting with the steam engine, there is a long tradition of regulating technologies in society through legal means. Regulators and policy makers engage in intensive discussions with the advent of new technologies. They are regularly looking for the best way to mitigate risks. While regulation, in its attempt to control technologies, often interferes with their development and uses, it also has functions that benefit technology development. It can create acceptance for a specific technology, set standards that might help the technology to be adopted by society and define criteria under which the use of this technology is acceptable. Therefore, regulation has a big impact on innovation and the extent to which technology thrives within a certain jurisdiction. After several countries formed strategies on AI and several proposals for ethical frameworks were developed, it was not surprising that regulatory discussions arose in many jurisdictions.

In order to talk about the regulation of AI, it is necessary to have a general understanding of what these terms mean. It is first important to note that AI is not one specific technology, but, in the words of Urs Gasser, “a set of technologies.” Klaus Mainzer has defined it as systems solving complex problems independently. This definition points to the generic nature of AI. It is rather a research question about machines being able to cope with complex problems. Computer scientists have developed many approaches to AI. In the 1990s, rule-based expert systems tried to create vast decision trees that would guide individual decisions. The most important approach within the field of AI today is machine learning. This is commonly defined as the process by which algorithms improve from data and experience. In contrast to rule-based programming, machine-learning models adapt themselves to their training data through certain mechanisms such as linear regression and backpropagation. To put it very simply, the systems learn gradually, using training data that contain a certain input and the respective result. Take, for example, the training of an image-recognition algorithm that is able to recognize cats. The training data would contain pictures with objects together with the information whether the picture contains a cat. If one picture is fed into the algorithm, it takes an initial shot, but then reviews its decision in light of the result and adapts the weights in the system gradually. The important difference from rule-based approaches is that explicit programming is not necessary. The system is designed in a way that it can improve by analyzing existing data. Machine learning has allowed significant breakthroughs, starting with image recognition. However, despite the great advances of machine learning, it will probably not be the last invention in the field of AI.

In order to understand what the regulation of AI means, we have to delve deeper into AI's general-purpose nature. As mentioned above, AI can be applied in different sectors, can improve existing technologies and gives also rise to innovations. In line with this, AI can be used for numerous actions in society, across a range of various tasks, including detection, knowledge management, communication, action and recommendation. It therefore, fuels several applications such as chatbots, intelligent assistants, and facial recognition devices or automated driving.



AI Applications [licence](#)
[info](#)

Due to its general-purpose nature, it is hard to pin AI down from a regulatory standpoint and to put it in a specific place having limited use cases and risks and opportunities. From a regulatory standpoint, AI can be addressed in different ways from a regional (local, national or international) and sectoral viewpoint.

Cross-border considerations for AI regulation are not as universal as in the case of other digital technologies like the internet. For example, AI applications such as automated cars do not necessarily entail a cross-border element. However, in many situations there is a manifest interest to create international rules, as in the case of lethal autonomous weapons systems, a very controversial use of AI by armed forces. Due to the international nature of conflicts, the discussion is mainly focused on the international plane. While the International Committee of the Red Cross has put this topic on the agenda several times, it is also being discussed by a Group of Governmental Experts under the framework of the Convention on Conventional Weapons. There are also international regulatory efforts, like the Council of Europe's [Ad Hoc Committee on Artificial Intelligence](#), which aims at producing a framework convention on AI. However, there are also regulatory efforts on the national plane, like a [law on algorithmic transparency in France](#). Even federal entities discuss legislative proposals on AI, as is currently occurring in the German federal entity Schleswig-Holstein. There can be more granular rules on the communal level. In the United States, there have been several bans or restrictions on facial recognition – for example, in [Oakland and San Francisco](#) in the state of California, or [Northampton, MA](#). AI regulation can also apply to various sectors, as European Union law shows. Sometimes AI is also regulated in the context of technologies in one sector. In the context of [medical devices](#), AI is just one of many regulated technologies. In contrast, the [EU's proposed Artificial Intelligence Act \(PAIA\)](#) is focused on a cross-sectoral regulation addressing all uses of AI.

Why is the regulation of AI important?

AI is not only a topic for the future; machine learning and other AI technologies are already in use today. The uptake of these technologies will increase in the coming years. Many countries have emphasized their willingness to [shape the development of AI through strategies](#). Some of these strategies contain regulative elements together with incentives and organizational measures. The analysis of those strategies reveals that they have different purposes, including supporting and furthering research, innovation, and adoption of the technology, but also mitigating risks.

Risk Mitigation

The discourse around the societal impacts of AI has revealed many potential negative impacts. A first set of discussions centered on discrimination, transparency and, accountability among others. One distinct feature of machine learning is the ability of a system to improve not through explicit human programming but using data. This further step of automation means that humans have no knowledge of the exact workings of an algorithm in the first place. AI, therefore, is opaque to its developers and users alike. This opacity runs counter to expectations of transparency (the extent to which a given AI system's inner workings are open to analysis) and explainability (the extent to which the information available for a certain decision can be understood by a specific stakeholder) of decisions in many areas of a democratic society. Legal norms require transparency and explainability in many cases. In public administration, for example, the rule of law requires decision makers to provide citizens with the rationale underlying their practices, to ensure that decisions and actions are based on solid grounds. The provision of rationale has to apply to algorithmic decisions as well.

There are numerous stories illustrating the importance of demystifying algorithms in terms of transparency and explainability. For example, a research [group from the Fraunhofer Heinrich-Hertz-Institute](#) found that an award-winning image-recognition algorithm had been making decisions on false premises. It could recognize horses in an image, but in fact, it was looking only for fences and similar structures as well as the copyright notice in the lower-left corner of the picture. This was because the algorithm was trained on data that contained such copyright notices and structures similar to fences. The algorithm was not actually fit to detect horses on a picture! This was only revealed after inquiries into the inner working of the algorithm. The practice of justifying decisions is often a question of degree. Full transparency might require documentation of the whole development of the design process and a thorough analysis of the algorithm. On the other end of the spectrum of transparency are general notions of the reasons for the algorithm to take action. Furthermore, there is a trade-off between transparency and the effectiveness of the process. Computer scientists [assume that transparency of AI systems could hamper their accuracy](#). From this perspective, there are several reasons for regulation to step in. First, regulation could improve existing practices and to ensure a reasonable standard of transparency and explainability. Second, regulation is the best tool for addressing existing trade-offs and striking the optimal balances between a full explanation and the overall effectiveness and efficiency of processes.

Transparency and explainability are also helpful in the consideration of additional questions concerning the performance of algorithms. Pertinent issues in this regard include equality, fairness and discrimination. In cases where algorithms are used to support or make decisions, the question arises as to whether these decisions are sufficiently fair

and whether various groups are treated equally. The use of AI in society has the potential to scale certain decisions up, thereby increasing the impacts of predispositions. One example is algorithms that sort applications for prospective employers. Their widespread use would mean that their predispositions concerning fairness become relevant for many organizations at the same time. A discriminatory algorithmic component would therefore have a much bigger impact. While equality and fairness are topics that are protected by human rights law and by specific legislation, the uptake of AI poses the question as to whether this legislation needs to be updated or specified for questions of algorithmic discrimination. Equality and fairness in algorithms depend on many aspects, some of which are exclusive to algorithms. One example is data governance and data curation. These relate to whether and how certain parts of the population are represented in the data. If an algorithm for skin-cancer detection is trained exclusively on the skin of persons of one color, it might not be able to detect skin cancer for other ethnicities. The focus on data governance also reveals potential trade-offs in the respective systems designs. From a standpoint of equality and fairness, it is important to collect as much data as possible, especially data on features that make up certain groups in categories such as ethnicity, gender, and religion. This data is necessary when checking for discriminatory effects. However, this violates principles of privacy that call for data minimization. Furthermore, it is possible that a strong emphasis on equality and fairness has a negative impact on the overall accuracy of the system. According to this trade-off, corrections of biases and discriminatory effects can hamper how well the systems function. These scenarios for trade-offs exemplify the intricacies of addressing questions of equality and fairness through regulation.

Another important topic under discussion is the accountability of AI. Does the adoption of AI lead to gaps in the responsibility for AI-induced actions? Beyond the mitigation of material risks, it is of crucial importance to establish who is responsible for malfunctions, the violation of rights, and damages incurred through AI. It is one of the major tasks of the legal regulation of technology to attribute responsibility in a clear-cut fashion to certain actors. In the case of AI, this is discussed in traditional and new legal constellations. Traditional questions point to the attribution of responsibility to certain actors including producers, sellers and users of technology. There are also different legal degrees to which certain users are responsible. These range from liability for every violation of the law to higher requirements regarding evidence or intent. Another proposal to enhance responsibility is to create a new legal status for machines. In legal terms, this would require that machines would have some form of personality, i.e. that they could be bearers of rights and duties. While a [report to the European Parliament](#) asked to define such a machine personality, this idea was later revoked.

While fairness, transparency and accountability can be seen as main points in the AI debate, an increasing number of issues point to additional risks. One question is sustainability of AI and in particular environmental sustainability. During the hype around AI, it is rarely highlighted what the energy costs of training, evaluating and using these systems are. This is just one example of the many risks a general-purpose technology might entail.

Furthering Technology

Legislation can have the effect of creating trust and acceptance. It is a democratic tool enabling intensive debates and discussions and the inclusion of several stakeholders. At the end of the legislative process, there is a text outlining the requirements for development and use of the technology as well as potential enforcement measures. Legal regulation thus sets a minimum standard for the technology. Legal regulation helps to “stabilize expectations” as Niklas Luhmann put it, thereby increasing the probability of engendering trust in the technology. This can create responsibility for the technology, but also define the limits of responsibility.

Regulating AI can also mean regulating certain questions differently. As new technologies allow for new features of the technology, an update might be necessary to do justice to new opportunities and risks. One example in this regard are risk management systems. They offer completely new possibilities for evaluating risks. Risk and dangers play an important role vis-a-vis many aspects of the law; however, the law assumes risk assessment by persons or groups of persons and is framed accordingly. Algorithmic risk assessment introduces new possibilities. The respective processes for quality management have to be designed differently, especially regarding quality management and review. Section 88 Subsection 5 of the [German Fiscal Code](#) sets out the following minimum requirements so that risk management systems must ensure that:

1. A sufficient number of cases are selected, on the basis of random selection, for comprehensive review by officials,
2. Officials review those cases sorted out as requiring review,
3. Officials are able to select cases for comprehensive review,
4. Regular reviews are conducted to determine whether risk management systems are fulfilling their objectives.

Regulation might also be needed to adapt processes in order to make it possible to realize the opportunities the technologies offer. Sometimes, regulation is needed to lessen the burdens of adopting a particular technology. In the case of AI, sandboxing – i.e. lifting regulatory requirements in certain cases – and experimentation clauses are suggested in different contexts, such as the [Norwegian AI strategy](#) or the EU’s PAIA. Art. 53 & 54 PAIA of the proposed PAIA lifts the obligation to process data destined only for specified purposes. If the requirements of Art. 54 PAIA are met, the data can be used to train algorithms irrespective of the purpose for which the data was collected. This is a massive lessening of the burdens.

Rebalancing Power Relations

Regulation is also important for recalibrating the power dimension of AI. AI is often perceived as a power factor in different relations as it allows new levels of knowledge and actions. Regarding knowledge, AI makes it possible to surveil situations ubiquitously and permanently, which can have a big impact on relationships. Employers are potentially able to surveil the performance of their employees and to detect certain behaviors. Regulation can address these changing capabilities by outlawing or limiting such surveillance or by introducing new rights for employees. The aforementioned lethal autonomous weapons systems could also create new military advantages that would change

the power relations between states. Again, regulation can help to balance these changes.

The timely regulation of AI can also be important from a standard-setting perspective. As we have seen above, it can substantially influence the way in which values are represented in technology. Swiftly enacted regulation in one jurisdiction can create substantive spillover effects to other jurisdictions. For example, the European Union's General Data Protection Regulation (GDPR) has a very broad extraterritorial application, as it applies, according to Art. 3 to "the offering of goods or services, irrespective of whether a payment of the data subject is required, to such data subjects in the Union; or the monitoring of their behavior as far as their behavior takes place within the Union." Given that the EU is an important market for service providers and data analysts alike, this has forced organizations outside the European Union to abide by its standards. Many companies have even implemented those standards generally, as it would have been more complicated and burdensome to distinguish between different jurisdictions. The GDPR has also substantially influenced legislative efforts in other countries like Brazil through the general regulatory ideas it expresses. This is why early approaches to regulation can influence other regulatory debates.

How did the current regulatory debate develop?

Aspects of AI were discussed very early in the process of digitalization. In 1978, the first French law on data protection included a rule on automated decision-making systems as well as on decision aides in public administration and in the judicial system. This inspired rules on automated decisions in many codes of data protection. The general debate on AI started when it became a topic in international politics in the light of advances in machine learning. A number of statements of heads of state led to continued attention: Russian President Vladimir Putin famously stated [at the Russian children's "Knowledge Day"](#) that "whoever becomes the leader in this sphere will rule the world and it would not be desirable for this monopoly to be concentrated in someone's hands." China has published its [strategy](#) to become the leading nation in AI by 2030. The UK's then Prime Minister Theresa May stated at the World Economic Forum that she was "establishing the UK as a world leader in Artificial Intelligence." Such statements led to more attention on AI by policy makers, which led to a significant rise in national AI [strategies in several countries](#). Regulation of AI was a part of many of those strategies. Policy discussions ensued in many countries. One of the processes receiving the most attention was the regulation of AI in the European Union. Member states issued a common declaration in early 2019; the Commission followed suit with an AI strategy that entailed regulation as one measure. When the new president of the Commission, Ursula von der Leyen, entered into office, she made the regulation of AI one of her top priorities, announcing the [following in her candidacy speech for elections to the European parliament](#):

In my first 100 days in office, I will put forward legislation for a coordinated European approach on the human and ethical implications of Artificial Intelligence. This should also look at how we can use big data for innovations that create wealth for our societies and our businesses.

The Commission followed up with a [whitepaper](#) and an ensuing stakeholder consultation. On April 21, 2021, the European Commission published its [proposal for an Artificial Intelligence Act \(PAIA\)](#), which set forth a general regulation of AI through a risk-based approach. The proposed regulation outlaws certain uses of AI completely, while requiring more detailed regulation for high-risk applications and certain transparency requirements for limited-risk applications. The proposal also included rules on enforcement including an obligation for member states to designate a national supervisory authority and the establishment of the European Artificial Intelligence Board. The proposal is a good example of the main topics that need to be discussed in the regulation of AI.

Controversial questions

While there are many specific AI-related issues such as liability in civil law and facial recognition, there are certain controversial questions of a general nature. Among them are the following issues.

Legal governance and innovation

One pertinent question is whether emerging technologies can and should be subjected to legal governance. A stereotypical view of regulation associates it with measures hampering creativity and innovation. As stated earlier, this is certainly not the only effect regulation can have. Therefore, the more interesting question is how we might imagine regulation that mitigates risks while empowering the realization of opportunities at the same time. Looking at the EU's PAIA, several such measures are apparent. The first is its risk-based approach that correlates regulation with the risk level of the technology. Systems that are not outlawed by Art. 5 PAIA and fall neither under the category of high-risk systems nor under the transparency obligations of Art. 52 are not addressed by PAIA. Other regulations are not that liberal. Take, for example, the GDPR. According to its Art. 2, section 1 it applies to the processing of personal data irrespective of the level of risk that is involved in the processing. There are also other provisions that are intended to stifle innovation, such as Arts. 53-54 PAIA, which have been mentioned above.

However, in all regulations, even in the PAIA, the challenge of balancing the goals of regulation with liberties to innovate always arises. The PAIA contains a detailed list of requirements for high-risk systems that reaches far into the development process and addresses many actors including sellers and users. These include the following:

- Setting up a risk management system (Art. 9)
- Measures concerning data governance (Art. 10)
- Technical documentation (Art. 11)
- Record keeping (Art. 12)
- Transparency and provision of information to users (Art. 13)
- Human oversight (Art. 14)
- Further material requirements concerning accuracy, robustness and cybersecurity (Art. 15)

These measures and their implementation mean a substantive load of work for providers and users of those systems.

The challenge of regulating AI is to regulate sufficiently while not imposing excessively heavy burdens on the actors on the ground.

Defining AI

An important question for the legal governance of AI is the very definition of AI. Given that AI is an umbrella term for a set of technologies that can solve complex problems independently, it is quite difficult to come up with a clear-cut definition that is also sufficiently pliable to accommodate future developments. In Art. 3, sec. (a) of the EU's AIA proposal, the actual definition can be changed through delegated legislation and is regulated in an annex to the regulation. However, its open wording shows that several problems lie ahead. For example Annex I includes the following:

a) Machine learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning; ...

(c) Statistical approaches, Bayesian estimation, search and optimization methods.

There are several controversies around this definition. First, it is not clear-cut, given that machine-learning approaches contain statistical approaches in many instances. Secondly, the terms “statistical approaches” and “search and optimization methods” are so broad that they could entail many pieces of software that nobody would normally associate with AI.

Outlawing certain uses of AI

One controversial question in processes of regulating AI is whether certain uses should be outlawed and forbidden in a general fashion. In this regard, it is very interesting to contrast the PAIA with the ethical framework designed by the [European Union's High Level Expert Group on Artificial Intelligence](#). This ethical framework does not designate forbidden uses of AI, but only mentions some critical concerns. In contrast, Art. 5 of the PAIA contains a number of uses that are to be forbidden under certain requirements such as ‘real-time’ remote biometric identification systems in publicly accessible spaces or AI systems that deploy subliminal techniques beyond a person's consciousness in order to materially distort their behavior. These regulatory red lines are subject to intensive discussions in different jurisdictions. While representatives of the industry generally oppose such prohibitions, activists and critics push to broaden them. However, regulation through legal governance must ultimately arrive at an actual conclusion. While red lines and prohibitions of certain uses of AI are hard to define, it seems that legal governance mechanisms are the most apt framework for defining those red lines.

Democratic AI governance

Another pertinent question that is frequently discussed is AI's impacts on decision-making and the consequences for democracy. As a concept, Techno-regulation highlights the increased capacity of technical systems to govern situations in society. AI's increased power in this regard is problematic. One way to address this issue is to introduce measures of human oversight. Art. 14 PAIA contains an obligation to design systems in a way that enables such oversight with the aim of “preventing or minimizing the risks to health, safety or fundamental rights.” It is significant that Art. 14 does not exclusively focus on AI systems, but also contains requirements under which humans can exercise such oversight. Art. 14 sec. 4 mentions particularly the competences and knowledge of the person overseeing the AI system. However, another aspect of democratic governance is the question who gets to build AI systems and whether the public ought to participate in these processes. General reflections on responsible research and innovation highlight this additional aspect of democracy and public participation. The EU's proposal is silent on these questions. Other regulatory efforts in more specific areas sometimes include participatory and democratic measures. One area where this might be particularly important is public administration. Several methods of deliberation can give citizens an active say when it comes to systems that take over important decisions in society.

Jurisdiction and international implications

As outlined above, the territorial applicability of regulation is of key importance. The question is how far such regulation should extend beyond the limits of the territory, especially when it comes to providers or users of the technology outside the jurisdiction that have effects within the jurisdiction. Such extension is problematic since it violates the sovereignty of other states and leaves the affected individuals without a say in legislation that affects them. A very restrictive approach to jurisdiction substantially limits the effectiveness of the regulation and allows for ways to circumvent it easily. Regulators have to take a stance on jurisdiction. An analysis of the EU's proposal shows the potential layers of the scope of such a regulation. Art. 2 sec. 1 (a) extends the scope to providers who place an AI system in the EU market, while Art. 2 sec. 1 (b) addresses users of AI systems within the Union. In contrast Art. 2 sec. 1 (c) focuses, on the effects of AI. It includes “providers and users of AI systems that are located in a third country, where the output produced by the system is used in the Union.” This clause is rather open. The reference to use of outputs of systems could potentially address any product that is at least partly produced with the help of an AI system. It will be interesting to see whether this open clause will be further restricted in the subsequent discussions concerning the proposal.

What's next?

Many of the proposals currently discussed will be passed into law through the legislative process. Amendments are to be expected. The PAIA is now being discussed within the [European Parliament and the Council of the European Union](#). The Council of Europe has not yet published a proposal. The first question that lies ahead is whether the current regulatory efforts will result in actual laws. Given the several legislative attempts and the pull to try to be the first to legislate, it can be expected that some of the proposals will turn into laws. The next question is how the approaches in jurisdictions compare to each other and whether there are conflicts in applying different laws. Such con-

flicts might prompt coordination mechanisms. There might be different forums for such a coordination. In the UN system, the Internet Governance Forum has taken charge of addressing governance issues of AI on a regular basis. Regional mechanisms such as the Council of Europe might help in building a common approach in regions. The current bilateral governance of international investment treaties might also include rules on AI in the future. Given that there are many controversial questions, it is probable that different approaches will emerge. This would then beg the question of which regulatory approach works best to mitigate risks and realize opportunities. The workings of the different regulatory approaches will have to be reviewed and evaluated in order to see how we can shape the future of AI through regulation.

Further Reading

Chander, S., & Jakubowska, E. (2021, April 28). "EU's AI Law Needs Major Changes to Prevent Discrimination and Mass Surveillance," *European Digital Rights (EDRI)*. <https://edri.org/our-work/eus-ai-law-needs-major-changes-to-prevent-dis...>

Cihon, P. (April 2019). *Standards for AI Governance: International Standards to Enable Global Coordination in AI Research & Development* [Technical Report]. https://www.fhi.ox.ac.uk/wp-content/uploads/Standards_FHI-Technical-Re...

Djeffal, C. (2019a). "AI, Democracy, and the Law." In A. Sudmann (Ed.), *The Democratization of Artificial Intelligence: Net Politics in the Era of Learning Algorithms* (pp. 255–284). Transcript.

Djeffal, C. (2019b). "Sustainable Development of Artificial Intelligence (SAID)." *Global Solutions Journal*(4), 186–192.

Djeffal, C. (2020). "The Normative Potential of the European Rule on Automated Decisions: A New Reading for Art. 22 GDPR." *Zeitschrift Für Ausländisches Öffentliches Recht Und Völkerrecht*, 81, 847–879.

"The EU Wants to Become the World's Super-Regulator in AI: The Brussels Effect," (2021, April 24). *The Economist*. <https://www-economist-com.emedien.ub.uni-muenchen.de/europe/2021/04/24/...>

Proposal for a Regulation Laying down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), COM(2021) 206 final (2021). <https://ec.europa.eu/newsroom/dae/items/709090>

High Level Expert Group on Artificial Intelligence. (2019). *Ethics guidelines for trustworthy AI*. <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-tr...>

Pagallo, U., & Barfield, W. (Eds.). (2018). *Research Handbook on the Law of Artificial Intelligence*. Edward Elgar Publishing.

Wischmeyer, T., & Rademacher, T. (Eds.). (2020). *Regulating Artificial Intelligence*. Springer.

The opinions expressed in this text are solely that of the author/s and do not necessarily reflect the views of the Heinrich Böll Foundation and/or of the Israel Public Policy Institute (IPPI), their staff, trustees and/or the organizations that support their work.

30 December 2021

by Christian Djeffal

© Heinrich-Böll-Stiftung e.V.

Schumannstraße 8
10117 Berlin
T +49 (30) 285 34-0
F +49 (30) 285 34-109
www.boell.de
info@boell.de