



Why “Artificial Intelligence” Should Not Be Regulated

DANIEL BRAUN, Department of Hightech-Business and Entrepreneurship, University of Twente, Enschede, Netherlands

Lawmakers all over the world have started to draft new regulations for Artificial Intelligence (AI). While the European Union is currently leading the way with its AI Act, many other legislators will follow and already positioned themselves with white papers and other publications. This commentary argues that “Artificial Intelligence”, including Generative AI, should not be used as a regulatory category. Not because there is no potential for harm from AI systems and not because AI systems should not be regulated, but because “Artificial Intelligence” is a vaguely defined label that is neither suitable nor necessary for comprehensive regulation of technological risks. Instead of regulating a particular set of approaches and algorithms, lawmakers should focus and double down on regulating high-risk applications of software, independent of whether it is labelled as AI or not.

CCS Concepts: • **Social and professional topics** → **Government technology policy**; • **Computing methodologies** → **Artificial intelligence**; • **Information systems** → **Decision support systems**.

Additional Key Words and Phrases: Artificial Intelligence, Generative AI, Regulation, AI Act

1 Introduction

68 years ago, the “Dartmouth Summer Research Project on Artificial Intelligence” gifted us with the term *Artificial Intelligence*. It has haunted us ever since. The word *intelligence* itself is ill-defined and whatever it describes is hard to measure in humans [11]. Not very surprisingly, prefixing it with *artificial* has done little to clear things up. For the best part of the last 68 years, the lack of a proper definition of AI has bothered, at most, a limited number of researchers. In recent years, the ambiguity of the term might have even been a blessing, at least to marketing departments that were able to market arbitrary products and services as AI, where AI seemed to be a synonym for technologically advanced or sometimes just “cool”. We have now entered a decade in which global lawmakers, for the first time, make serious attempts to draft comprehensive regulatory frameworks for AI. A task that comes with many challenges. One of them is the necessity for a clear definition of what is to be regulated, i.e., of what AI is. Looking at attempts made so far, outlined in Section 4, it is clear that among lawmakers, there is also no universally accepted definition of AI. This commentary challenges the very idea of regulating “Artificial Intelligence”. It will show that definitions used by lawmakers often only cover a subset of what could reasonably be seen as AI, by focusing on machine learning (ML). It will argue that any meaningful definition of AI includes such a broad and diverse range of technologies that the characteristics shared by such technologies are so small that they do not form a useful category for targeted regulations and are virtually indistinguishable from non-AI software systems.

Nevertheless, the widespread and mostly unregulated use of AI and other systems in decision-making and other tasks poses a threat to individuals and our society at large. Instead of focusing on certain technologies, like AI, we argue that lawmakers should focus on regulating applications and fairness and transparency within processes irrespective of technology. Whether a social score is calculated based on statistically derived implicit

Authors’ Contact Information: Daniel Braun, Department of Hightech-Business and Entrepreneurship, University of Twente, Enschede, Overijssel, Netherlands; e-mail: d.braun@utwente.nl.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

ACM 2639-0175/2024/9-ART

<https://doi.org/10.1145/3696010>

rules or a set of human-crafted explicit rules (or even manually by a human) does not change the potential negative impact it can have. Therefore, lawmakers should focus on harmful applications of software in general rather than technologies to implement these applications.

This text will particularly focus on the European Union’s (EU) AI Act, which is widely believed to be potentially influential in future global AI legislation [20, 33, 34, 41], even by its critics [3, 24]. While the AI Act explicitly formulates the goal to take a risk-based approach and regulate harmful applications rather than technology, the commentary will argue that an *AI act*, i.e. a law to regulate a specific set of technologies called AI, can never truly achieve this goal. Moreover, the AI Act falls short of its own goal by explicitly excluding, e.g., rule-based systems and regulating what the AI Act calls “general-purpose AI models” independent of any concrete applications and their risk.

2 Related Work

It is widely acknowledged that one of the major challenges that regulators face today is the rapid advancement of technology, which makes the drafting of “future proof” legislation more challenging than ever [15, 28, 40]. While we just now started to see the first comprehensive regulatory frameworks that attempt to regulate AI, lawmakers and scholars have been thinking and working on regulating certain aspects of AI for a long time.

2.1 Challenges in Regulating AI

In addition to the general challenges that come with regulating a fast-developing field, the absence of a comprehensive definition of what AI actually is has been identified as a challenge for regulators before.

In 2019, for example, Buiten [9] discussed possible ways of regulating AI and identified as one of the problems that the “various definitions of AI used in the literature may be helpful to understand AI, but are unsuitable as a basis for new laws” [9]. Buiten argues for a more concise regulatory target, mostly aiming at machine learning. Smuha [40] in 2021 similarly noticed that “as of today, no common definition for AI exists, though a number of attempts [...] have been made in this regard” and that “the absence of a commonly agreed definition poses certain obstacles” in the regulatory process [40]. Hacker [25] identified the same problem “basically every technical book on AI uses its own definition so that even computer scientists have not yet been able to agree on a uniform concept of AI” and suggests “to speak, instead of AI, of machine learning (ML) techniques, which are much more clearly defined” [25].

While the absence of a common definition of AI and the implications on the regulatory process have been discussed before, the common theme in previous work was, also specifically with regard to the AI Act [23], that definitions of AI are problematic because they are too broad and include too many systems. Scholars have therefore suggested, that regulation should focus on ML instead of the more generic AI. This commentary, on the other hand, will argue that, if anything, the AI definitions used by lawmakers are often too narrow, because they exclude, for example, rule-based systems. A truly technology-neutral regulation that focuses on the risks of applications rather than technology should not make such a distinction.

2.2 AI Act Critique

The AI Act, on which this commentary will focus in particular, in its different drafts, has been critically analysed by many scholars. One of the most critiqued aspects is that the AI Act itself does not grant any rights to those impacted by the systems it aims to regulate and hardly considers them at all (see e.g. [14, 44]). This critique is connected with the assessment that the AI Act, while at core a product safety regulation, mixes it with aspects of fundamental rights protection without sufficiently accounting for the differences in both areas and ultimately falls short in both of them [2].

While the risk-based approach the AI Act aims to take is widely seen as the right direction and sometimes even described as a “necessity” for regulating AI [49], the concrete implementation has been subject to different critiques. Edwards [14], for example, concluded that “The alleged ‘risk-based’ nature of the Act is illusory and arbitrary” because the act lacks reviewable criteria for the risk assessment, which they argue makes the AI Act “unacceptably arbitrary, denying justiciability and lacking futureproofing” [14]. Similarly, Paul [36] argues that “The [EU’s] AI regulation introduces neither comprehensive risk analysis methodology nor does it create robust independent risk assessment units” [36]. Fraser and Bello y Villarino criticise the AI Act for taking an “AFAP [(as far as possible)] approach to risk acceptability” in which “a vaguely defined ‘state of the art’ [...] dictate[s] how much risk society should bear from high-risk AI systems” [22].

Many scholars believe that the AI Act will potentially influence global AI regulation (see e.g. [3, 20, 21, 24, 33, 34, 41]). Others believe that the AI Act will not have the widely cited “Brussels effect” [6], but rather start a “race to the bottom”, in which global lawmakers try to be attractive for businesses by putting less regulations on the usage of AI [37, 45].

3 “Artificial Intelligence” in Scholarly Literature

From its very beginning, the meaning of the term “Artificial Intelligence” was connected to human abilities. The organisers of the Dartmouth project wanted it to be about “how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans”. Even before the term was coined, scientists were thinking about what would make machines intelligent. Alan Turing, famously, discussed the issue in his 1950 article “Computing Machinery and Intelligence” [42]. He suggested the “imitation game” as the ultimate test for machine intelligence. An adapted version of which today is known as the Turing test. Both, the original and the adapted version, at the core define a machine as intelligent if it can deceive a human into believing the machine is human. If we learned anything from two decades of the Loebner prize, a competition to develop machines that can pass the Turing test, it is probably that the ability to deceive a human in such a way is not necessarily correlated to our everyday understanding of *intelligence*. On the contrary, those systems that have been more successful in the competition were usually systems that did not indulge in attempts to *understand* what was said, but rather used clever linguistic re-arrangement techniques to keep a conversation, just like ELIZA did in 1966 [46]. Thirty decades later, in his 1980 book “Principles of Artificial Intelligence”, Nilsson [35] wrote “Many human mental activities such as writing computer programs, doing mathematics, engaging in commonsense reasoning, understanding language, and even driving an automobile are said to demand ‘intelligence’. Over the past few decades, several computer systems have been built that can perform tasks such as these. [...] We might say that such systems possess some degree of *artificial intelligence*” [35].

More than 70 years after Turing’s thought experiment, we are not much closer to a widely accepted definition of AI in science. In their standard reference “Artificial Intelligence: A Modern Approach” Russell and Norvig [38] do not provide a single definition of AI, but rather two dimensions alongside which different definitions of AI have been given by researchers throughout the years. In the first dimension, AI is defined as either *human* or *rational* and in the second dimension as either *thought* or *behaviour*-based. That leaves us with four possible definitions of AI as either systems that act human (like Turing thought), think like humans, act rational, or think rational (see Figure 1). [38, p. 19-22] All four ways to define AI include some degree of anthropomorphisation and are ultimately based on what we define as human-like or rational. Although evidently not the same, rationality in the end is something that we almost exclusively attribute to humans.

A trend that cannot only be observed in the scientific usage of the term AI, but also in everyday usage and, as will be shown, in legislation, is to use it interchangeably with ML. Henman [26], to name just one example, writes “AI is typically used to refer to (systems developed with) machine learning algorithms” [26]. As others have pointed out before (see e.g. [29, 52]), AI and ML are by no means synonyms. ML is a sub-field (or sub-set)

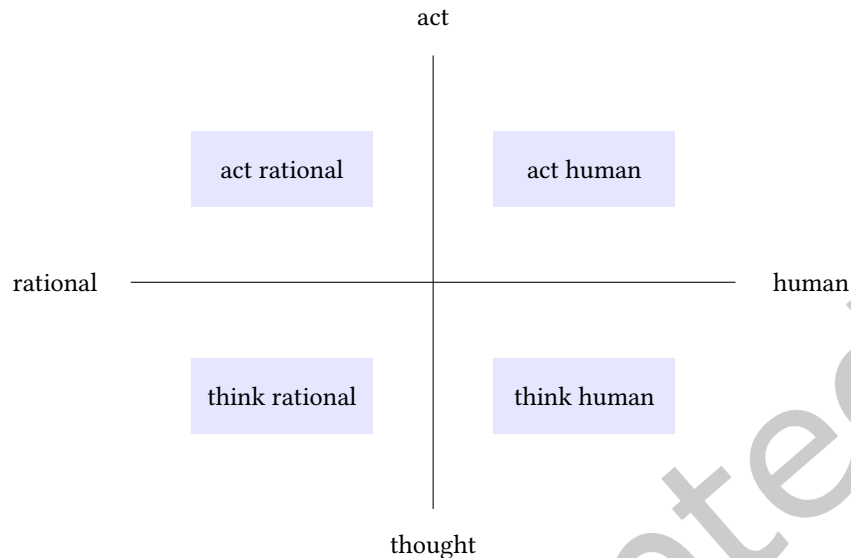


Fig. 1. Different dimensions of defining AI according to Russel and Norvig [38]

of AI, one of multiple. One important branch of AI that is particularly excluded when equating it with ML are so-called *expert systems*. Expert systems are *knowledge-based systems* in which a symbolic representation of human knowledge is applied in order to offer solutions to specific problems in a given domain [31, p. 1]. Expert systems usually consist of two components, a knowledge base which contains the domain knowledge that is usually encoded in some type of formalism and an inference engine which applies logical rules to the knowledge base in order to deduce or infer new information from the available knowledge [31, p. 7]. Some of the first commercially successful AI systems have been rule-based expert systems [38, p. 40-42].

Today, many reject the idea that such systems should be classified as AI and would rather classify them as “just a bunch of if-statements”. Even in the scholarly literature, there is a trend towards granting the label AI only to the latest and most exciting technologies. In his widely cited book “Introduction to Artificial Intelligence”, Ertel [16] suggests that the AI definition of Elaine Rich “elegantly” solves the dilemma of other definitions. It reads: “Artificial Intelligence is the study of how to make computers do things at which, at the moment, people are better.” (as cited in [16, p. 2]). Ultimately that would mean that as an approach improves and achieves “human-level” performance, it would stop being AI.

4 Regulation of AI

AI has been a subject of political debates and regulation for years. The EU’s General Data Protection Regulation (GDPR), for example, states in Article 22 “The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.” Effectively forbidding purely AI-based decisions in certain contexts. While regulating how AI can be used is not exactly new, what is new about the current efforts to regulate AI in Europe, the USA, Africa, China, and other jurisdictions is the attempt to create what is sometimes called a “comprehensive” or “horizontal” regulatory scheme, i.e. laws that regulate all, or at least most, aspects of AI technologies. For such regulations, a definition of “Artificial Intelligence” is indispensable. Lawmakers did not

seem to be impressed by scientific attempts to define AI, at least they chose different approaches for their own definitions.

4.1 USA

Some simply chose to ignore the issue by using AI without defining what it is. The “Blueprint for an AI Bill of Rights” by the White House, for example, provides several definitions, for concepts such as algorithmic discrimination, sensitive data, and surveillance technology, but abstains from defining AI, the very concept it wants to regulate. In the “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence”, the white house used the AI definition given in 15 U.S.C. 9401(3) “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.” This definition is anthropomorphising AI by attributing the ability to “perceive” to AI systems. While one could argue that complex computer vision systems, e.g. in cars, can “perceive” their environment to a certain degree, it seems far-fetched to claim more simple systems, like automated hiring systems (see Section 5.2), have perceptive abilities and would therefore fall under this definition. In general, the definition is clearly targeted at ML systems, because it only considers systems that “abstract such perceptions into models through analysis in an automated manner”.

4.2 Africa

In a significant step towards a comprehensive AI regulation in the African Union (AU), the African Union Development Agency (AUDA), released a white paper on AI regulation [1]. The white paper offers the following definition: “Artificial intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to process information and function similarly to human beings. In mimicking humans’ actions, AI-based systems are exhibiting traits that are associated with a human mind such as learning and problem-solving. AI is characterised by its ability to rationalize and take actions that have the best probability of accomplishing a specific goal. Machine learning (ML), as a subset of AI, deals with computer programmes that can automatically learn from and adapt to new data without being assisted by humans. On the other hand, deep learning, as a subset of AI, involves techniques that can enable automatic learning by absorbing quantities of unstructured data such as text, images, and videos.” [1] More than other definitions used by legislators, this definition resembles the scientific definitions as described by Russell and Norvig [38] by resembling the “act human” (“In mimicking humans’ actions”) and “think human” (“machines that are programmed to process information and function similarly to human beings”) perspective.

4.3 Europe

One of the currently most discussed pieces of AI legislation is the EU’s AI Act. Its process of formation exemplifies the issue of finding a legal definition for what AI is. The legislative process at the European level involves three entities: The European Commission, which has the right of initiative, and the Council of the EU and the European Parliament as legislators. In the legislative process leading to the final version of the AI Act, all three bodies made different proposals, notably not just on how to regulate but also on how to define AI. The initial draft of the European Commission defined AI as “software that is developed with one or more of the techniques and approaches listed in Annex I and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with.” [17]

The rather extensive, and at times quite specific, list of techniques and approaches in the annex included:

- (1) “Machine learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning;
- (2) Logic- and knowledge-based approaches, including knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning and expert systems;
- (3) Statistical approaches, Bayesian estimation, search and optimization methods.” [17]

This definition proposed by the Commission is based on the used technology as well as its output. However, both aspects are defined so broadly, that some experts said it is hard to think of a software that would not fall within this definition [10]. If it had been indeed the intention of the commission to regulate all kinds of software, it would seem unnecessary to use and define the term AI in the first place. The fact that there is an explicit list of techniques that have to be used in order for software to be considered AI strongly indicates that the Commission had other software in mind that is not seen as AI. The phrasing “logic-based approaches” alone would basically describe every software possible and the Commission either had a more restrictive interpretation in mind or did not sufficiently consider the implications.

Sharing the sentiment of some experts that the definition proposed by the Commission might be too broad, the Council of the EU proposed a different definition which, in their perspective, is more narrow, defining AI as:

“a system that is designed to operate with elements of autonomy and that, based on machine and/or human-provided data and inputs, infers how to achieve a given set of objectives using machine learning and/or logic- and knowledge based approaches, and produces system-generated outputs such as content (generative AI systems), predictions, recommendations or decisions, influencing the environments with which the AI system interacts” [12]

In addition to removing the explicit mentions of individual technologies and “statistical approaches” altogether, this definition introduces “elements of autonomy” as a property of AI systems. An interesting thought that excludes “simpler” software from the AI definition. Arguably, it also introduces even more uncertainty and vagueness about what an AI system is, because it raises the questions of what constitutes “elements of autonomy”.

Finally, the European Parliament used an adapted version of the definition used by the Organisation for Economic Co-operation and Development (OECD) prior to November 2023. Arguably, this presents a more concise version of what the Council wanted to say:

“artificial intelligence system (AI system) means a machine-based system that is designed to operate with varying levels of autonomy and that can, for explicit or implicit objectives, generate outputs such as predictions, recommendations, or decisions, that influence physical or virtual environments;” [18]

Figure 2 shows a visualisation of the concepts mentioned in the three definitions. While the definition of the European Parliament is the shortest, it might well be the most far-reaching one, depending on which notion of *autonomy* one is using. Notably, all three definitions try to avoid anthropomorphising AI or defining it based on human traits.

In November 2023, partially in response to the drafting of the AI Act, the OECD updated its definition of AI (systems). The new version of the definition reads as follows (additions in **bold**, removed parts ~~striktthrough~~):

“An AI system is a machine-based system that can, for a given set of human-defined **explicit or implicit** objectives, **infers, from the input it receives, how to generate outputs such as** makes predictions, **content**, recommendations, or decisions **that can influence** physical real or virtual environments. **Different** AI systems are designed to operate with varying **in their** levels of autonomy **and adaptiveness after deployment.**” [39]

In the final version of the AI Act, which was approved by the parliament on the 13th of March 2024, the new definition of the OECD was adopted:

“‘AI system’ means a machine-based system designed to operate with varying levels of autonomy, that may exhibit adaptiveness after deployment and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments” [19]

Removing all the optional parts (“may”, “such as”, “varying”) the definition boils down to a system that infers how to produce output based on the input it receives. Since (almost) every computer program receives input processes it and produces output (the so-called input–process–output (IPO) model), the question of how far-reaching this definition is depends largely on the interpretation of the word *infers*¹. As described in Section 3, expert systems have a component called “inference engine”, which *infers* based on the knowledge base. When a trained ML model is used to produce an output based on an input, the process is regularly referred to as *inference* [50]. So from a scientific perspective, both ML and rule-based approaches can clearly infer. However, it is not necessarily clear when “inference” starts. Is a calculator application inferring how to generate the output for a given input? In the end, it will be up to courts of law to interpret how far-reaching the definition of AI in the AI Act really is. However, the act offers an interpretation aid: The preamble of legislative acts in the EU regularly contains a so-called recital. Recitals are explanatory texts that “set out the reasons for the contents of the enacting terms (i.e. the articles) of an act” [43]. While the recital is not legally binding, it is often used as an interpretation aid that can (but does not have to be) used by courts when interpreting the law. The proliferation of the usage of recitals in EU acts is often criticised and linked to the poor quality of legislation at the EU level [13]. The recital of the AI Act explains under No. 12 the reasons for the chosen definition of AI systems. According to the text, among other points, the definition “should be based on key characteristics of AI systems that distinguish it from simpler traditional software systems or programming approaches and should not cover systems that are based on the rules defined solely by natural persons to automatically execute operations” [19]. While this reads very clearly like a specific exclusion of systems that use explicitly modelled rules that are based on expert knowledge, like expert systems, the text continues: “A key characteristic of AI systems is their capability to infer. [...] The techniques that enable inference while building an AI system include machine learning approaches that learn from data how to achieve certain objectives, and logic- and knowledge-based approaches that infer from encoded knowledge or symbolic representation of the task to be solved.” [19] At least to the author of this text, explicitly excluding “systems that are based on the rules defined solely by natural persons to automatically execute operations” and at the same time explicitly including “logic- and knowledge-based approaches that infer from encoded knowledge or symbolic representation” is an oxymoron.

5 Why We Do Not Need to Regulate “Artificial Intelligence”

If all definitions of AI that are currently used by lawmakers have their flaws, one could think that we just have to find a better definition of AI. However, instead, we should ask the more fundamental question of whether “Artificial Intelligence” is a meaningful categorisation for the regulation of technology. Lawmakers regulate technologies to prevent that they harm citizens, businesses, or society at large. Therefore, how heavily a technology is regulated usually depends on the potential harm it can cause and the potential benefits it can bring. This section will argue that using AI as a regulatory category is misguided because the label is not indicative of the amount or type of danger a system is posing. Focusing regulation on AI, rather than risk, can lead to over-regulation of AI systems that do not pose harm and under-regulation of systems that pose harm but do not fit the chosen definition of AI.

¹and whether “varying levels of autonomy” also includes little to no autonomy.

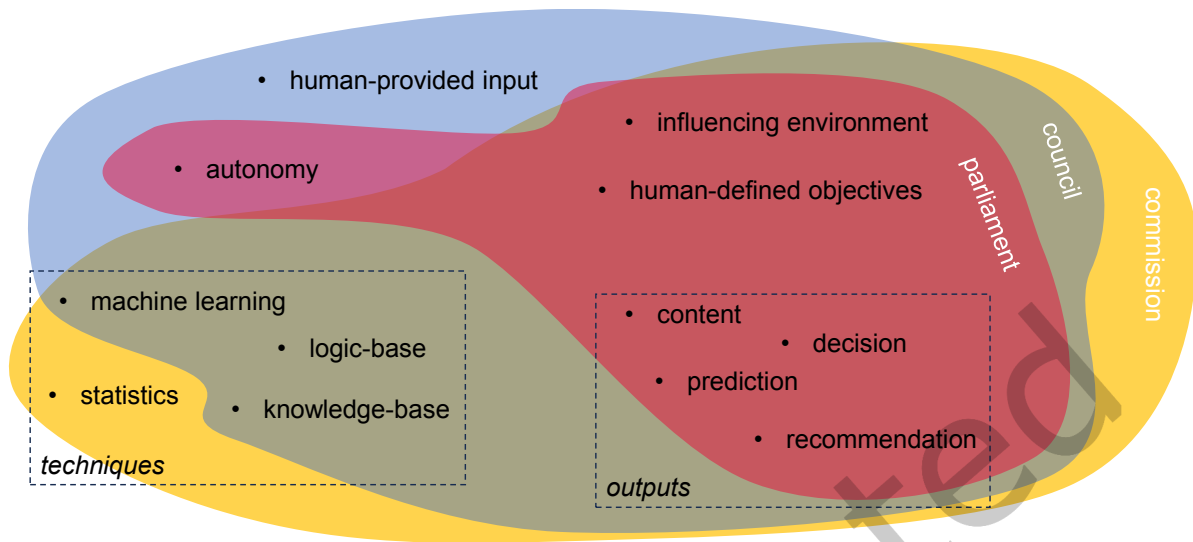


Fig. 2. Concepts mentioned in the definitions of AI used in the initial drafts of the AI Act by the Parliament, Council, and Commission of the European Union

5.1 Risk-based Approach within the AI Act

In acknowledgement of the challenges in regulating a fast-developing technology like AI (see Section 2.1), at its core, the AI Act is based on assessing the risk that certain applications of AI pose and then regulating them based on that risk rather than the underlying technology. A similar approach was already taken as part of the GDPR.

Based on the risk assessment, systems are more or less strictly regulated ranging from a complete prohibition for systems with “unacceptable” risks, like social scoring systems, to no additional regulation for minimal risk systems, like AI in computer games. Often, this risk-based approach is visualised with a pyramid similar to what can be seen in Figure 3. A fact that is often overlooked by such graphical representation is that the AI Act also defines a category of “general-purpose AI models” that is regulated independently of any application in a fashion that could be seen as somewhere between high- and limited-risk models. The regulation of general-purpose models will be discussed in more detail in Section 5.3.

The AI Act has come a long way in actually achieving a risk-focused and technology-neutral approach, compared to the the initial proposal of the Commission to list a set of technologies in the Appendix of the act. Neither writing an if-statement nor training a statistical model is inherently dangerous. However, they can be used in settings where their usage can cause harm to people, in one way or another. The AI Act rightly acknowledges this and seeks to regulate such usage. But why only for (however defined) AI? An AI act in itself can never truly be based on application risks rather than technology, because it assumes a priori that there is a set of technologies (AI) that is more risky and therefore needs special regulation. We will use an example of a high-risk application to illustrate why instead of using AI as a regulatory target, lawmakers should double down on the idea of assessing application risks and regulate *software*, rather than AI, systems that are applied in (high-)risk areas.

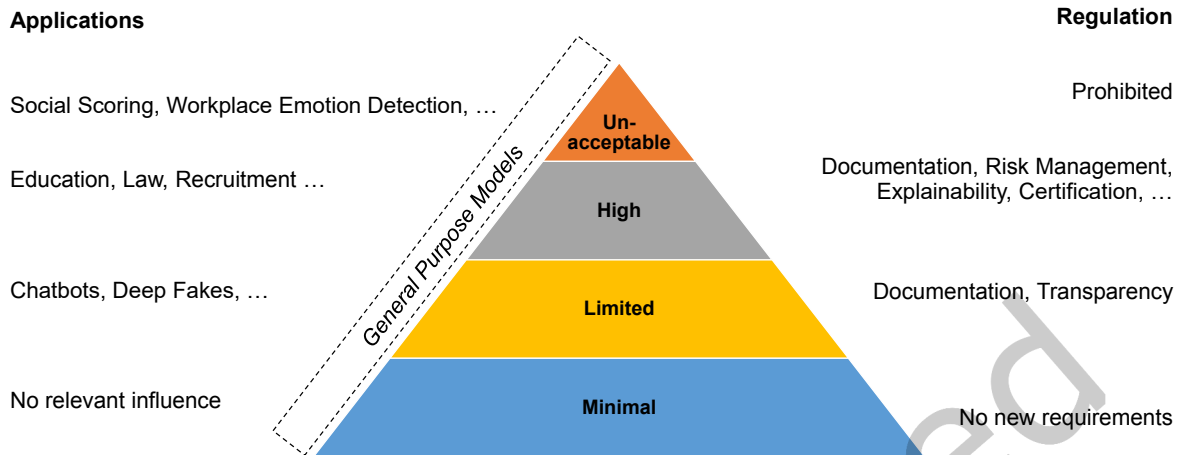


Fig. 3. Risk-based approach with the AI Act

5.2 High-risk Application: Hiring Systems

AI systems are used in the hiring process, e.g. to more efficiently screen applications or proactively find and approach suitable candidates [30]. In the past, such tools have proven to show different biases, e.g. with regard to gender [47]. Such systems are classified as high-risk systems in accordance with Annex III of the AI Act: “AI systems intended to be used for the recruitment or selection of natural persons, in particular to place targeted job advertisements, to analyse and filter job applications, and to evaluate candidates”. Among other things that means that for such systems the data used to train them (if any was used) should be “sufficiently representative” (Art. 10 No. 3), the system and the decisions it makes should be well documented (Art. 11), and effective human-oversight must be possible (Art. 14) [19].

Let’s assume three systems (A, B, and C) designed to support the process of selecting which applicant to invite for an interview. For the sake of simplicity, we assume that applicants fill in a structured recruiting form, which is the case in many big organisations, including their degrees and grades.

- (a) System A uses a random forest model that is trained on historical data. The data consists of filled-in recruiting forms from the past that are annotated with information on whether a candidate was invited for an interview or not.
- (b) System B is an expert system. The recruiting manager can manually define criteria for which to filter. Additionally, the system uses a knowledge base that is, among other things, able to convert between different grade systems to account for the fact that a very good grade for a Bachelor’s degree from, for example, Germany ranges between 1.0 and 1.3, while in the Netherlands it ranges from 10 to 9.
- (c) System C is a simple low-code system in which the user can define straightforward filters, e.g. whether the grade of an applicant is ≥ 9 .

Without any doubt, system A would be an AI system in the sense of all common AI definitions and the AI Act and therefore subject to the regulations of high-risk systems. Notwithstanding the ambiguity described in Section 4.3, system B is a classical expert system and therefore an AI system in the sense of the AI Act and most other definitions that do not equal AI and ML. System c, however, would not be seen as an AI (system) by most of the common scientific definitions and also not by the AI Act.

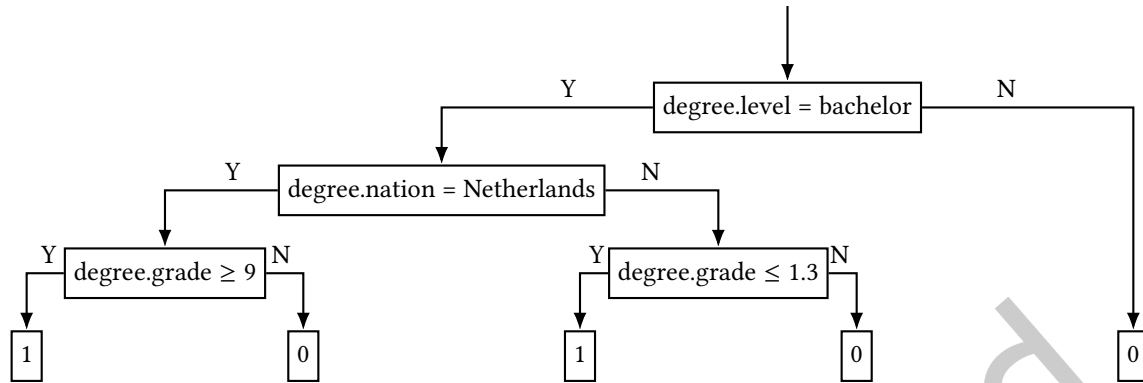


Fig. 4. Possible decision tree learned by system A (i.e. a system that uses a random forest model that is trained on historical data)

Let's assume all three systems are used by a "bias-free" Dutch company that has historically based its decisions of whom to invite solely on achieved grades and wants to continue to do that. If the company, so far, only received applications from people with degrees from Dutch and German universities, one of the decision trees learned by system A could potentially look like the tree shown in Figure 4. The tree would also describe how system B derives its decision, with potentially additional subbranches for grades from other countries depending on the extensiveness of the knowledge base. System C would effectively mimic the sub-tree on the bottom left in Figure 4. All three systems would therefore apply very similar criteria and all three systems would potentially discriminate against applicants with degrees from countries that use different grading schemes (e.g. from Switzerland, where 6 is the best possible grade).

While systems of all these types exist in the real world, this is of course a very simplistic scenario. Decisions made by such systems could, independent of the AI Act, be considered discriminatory². However, the example shows that independent of how complex a system is, how it ultimately makes a decision based on the input can be very similar or even the same. More importantly, the consequences it has are also the same. Of course, there are systems which introduce additional risks, e.g. neural networks that are, unlike random forests, not inherently explainable and therefore pose the risk of potential biases going undetected for longer.

However, the AI Act is neither an ML nor a deep learning act. It also includes expert systems like system B. From a technical and a risk perspective, such systems are so close to systems like C that it neither seems feasible nor necessary or useful to differentiate them with a definition of AI. If a system gives recommendations on hiring decisions, we should demand transparency, explainability, and proper documentation of its design, independent of what technology is used by the system. Because the harm a system can cause is not mainly determined by the technology it is using, the regulatory guidelines should not be either. Instead of nitpicking on AI definitions, (European) lawmakers should double down on their very own idea of regulating application risks rather than technology and make the AI Act a Software Systems Act and regulate all software systems that have a significant impact on our lives and not just AI. For developers of more simple systems, it would be relatively easy to comply with such an act and if a company is providing products that make such potentially life-changing decisions it does seem reasonable to demand those things from them, even if their system is not based on AI.

²Proving that, however, would be very difficult if the system is not subject to regulations like the AI Act that demand to log the decisions made and how they were made.




Fig. 5. Development of maximum amount of gigaflops available in different Nvidia GeForce desktop GPU series over time (based on data from [48])

5.3 “General-Purpose” “Artificial Intelligence”

It is widely discussed how the AI Act takes a “risk-based approach”. And that has very much been, and to a large extent still is, the core idea of the act. However, the proliferation of generative AI models, like the GPT models of OpenAI, has caused the EU to diverge from the original idea of just regulating AI systems, rather than the underlying AI models. Lawmakers deemed some models so risky that they deviated from two core principles of the AI Act to regulate them: For all other cases, the act never regulates an AI model itself, but rather the system that uses the model. And, most obviously, while AI systems are regulated based on the risk of their application, general-purpose models³, as they are called within the AI Act, are regulated independent of any application.

How far the regulation of general-purpose models deviates from the rest of the act becomes very clear in Art. 51 No. 2, which defines general-purpose models that are considered to have systemic risk: “A general-purpose AI model shall be presumed to have high impact capabilities [...] when the cumulative amount of computation used for its training measured in FLOPs is greater than 10^{25} ”. In the recitals (no. 111), this is justified by the fact that “According to the state of the art at the time of entry into force of this Regulation, the cumulative amount of compute used for the training of the general-purpose AI model measured in floating point operations (‘FLOPs’) is one of the relevant approximations for model capabilities.” That means that, for example, GPT-3 would not be considered to have high impact capabilities because the total train compute was 3.14×10^{23} [7]. While the use of more computational power can indeed often lead to better models, if, for example, the dataset used for the training is bad, the invested computational power becomes almost meaningless. Especially in settings with limited data, additional training of Large Language Models (LLMs), like GPT-3, which needs additional computing power, can even lead to decreasing results [32, 51].

Even if we accept computational power as an (imperfect) predictor for the capabilities of a model: If the general idea of the AI Act is to be future-proofed with regard to technological developments, it is an odd choice to add such a specific number that is so deeply rooted in the current state of the technology and so quickly changing. According to a 2022 report by Hobbhahn and Besiroglu [27], the FLOPs per Dollar double every 2.5 years. Figure 5 shows how rapid the development in available computing power has been in recent years even in the desktop segment.

³The AI Act defines a general-purpose AI model as “an AI model [...] trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks” (Art. 3, No. 63)

LLMs can be perceived as “general-purpose AI”, because they seemingly have the ability to perform a wider variety of tasks. However, in their training, they are trained for one specific purpose: to predict the most likely series of tokens given a specific input often called a prompt. Whether LLMs indeed have the ability to genuinely solve different tasks and maybe even show “sparks of artificial general intelligence” [8] or whether they are just “stochastic parrots” [5] that are simply trained on such a wide variety of data that they are able to produce useful looking output in so many cases that it inspires our imagination and leads us to attribute them with human traits [4, 5], is subject to fierce scientific debate. No matter which side of the aisle one is on, an LLM that is trained but never applied cannot have any negative impact beyond the environmental impact caused by the training of the model⁴.

LLMs can be used in many ways, including in a very restricted way on clearly defined tasks. LLMs can, for example, be used to generate headlines for news articles. In theory, they can even be used for such restricted tasks as generating passwords. It is not obvious why, if applied in such a restricted way, they should be regulated differently than other, classical, ML models. In practice, they are of course often applied for more complex tasks that bear higher risks. However, if an LLM is for example applied in a hiring system, the system will anyway be regulated based on the application risk, like other AI systems, which leads to a more strict regulation. One notable difference is that in the case of “normal” AI models, the responsibility to fulfil the regulations of the AI Act lies with the provider of the system that uses the model, which is not necessarily the same as the provider of the model. For general-purpose models, the provider of the model is bearing direct responsibility. In practice, this could well be a distinction without a difference. In order for their models to have a commercial value after the AI Act comes into effect, providers of general-purpose AI models would have to guarantee their customers compliance of the model with the act anyway. Otherwise, their models could not be used in commercial applications with risk anymore.

Together with the fact that separate regulation of general-purpose models was not part of the original draft, it seems like rather than being about the underlying models itself, the regulation is a response to applications that make these models directly available, like ChatGPT. ChatGPT is essentially a frontend provided by OpenAI that connects the user (presumably through some safety layers) with their GPT-3.5 model. ChatGPT itself is not built with any specific application in mind other than creating a chat-style interaction. Which rules would apply to applications like ChatGPT, if there is no specific regulation of general-purpose models? While this is an important question (that is to be discussed by legal scholars), it is not in any way unique to general-purpose models. There are, for example, commercial tools available that perform sentiment analysis. These tools are based on special-purpose, often relatively simple, models that are trained to detect the sentiment of a text. They are, for example, used by companies to analyse the reviews their products receive. An application that arguably bears very little risk since it does not have any relevant influence on the authors of the reviews that are analysed. However, the very same tools are equally useful to analyse annual reviews written by managers to make decisions about bonus payments. In that sense, general-purpose models are no different than other models in the fact that they can be used in different risk contexts. Instead of diverting from the very idea of the AI Act by introducing the category of general-purpose models, the AI Act should have addressed the problem that models and systems can be applied in different contexts within its risk-based framework. With the current regulation, it is difficult not to feel like the separate regulation of general-purpose models might not have been purely based on a regulatory need. After all, an AI Act that does not specifically address generative AI might have been seen as not comprehensive to the public, even if its regulations would still apply to them.

⁴While this impact should not be neglected, it is not something that is considered within the AI Act.

6 Conclusion

Letting machines make decisions can pose dangers to their users, to the people influenced by their decisions, and to society at large. As systems making decisions related to important aspects of our lives will become more and more common, lawmakers do well to regulate these. The EU’s AI Act is an important and laudable step towards comprehensive regulation of such technology. Compared to other proposed regulations, the AI Act, in theory, follows the right idea of regulating technology based on application risks, rather than the underlying algorithm. In practice, however, it seems like, in the end, the urge for pioneering dedicated AI regulation and responding to the public debates about generative AI was bigger than the legislative will for comprehensive risk regulation, ultimately watering down the idea of risk-based regulation.

This comment argued that legislators should stop trying to define “Artificial Intelligence”. Not only because, as of today, we simply do not have an accepted definition of AI, but because any definition of AI will always be based on a differentiation from other software. A differentiation that is, at least if one believes in the risk-based approach, not sensible. The real-world implications a software can have do not depend on the underlying technology and ultimately AI is nothing but an umbrella of, however defined, technologies. Therefore, legislators should stop trying to come up with new definitions of AI and focus on addressing the risks that different applications of software can have.

Scholarly debates about new regulations are naturally mostly led by legal scholars. As discussed in Section 2.2, there is a lively debate about the quality of the AI Act and other AI regulations. Despite its best efforts, even the AI Act, which was designed to focus on risks rather than technology, is deeply technical. It talks about context windows, accuracy levels, adversarial attacks, FLOPs, and makes implicit assumptions about the risk of different approaches. This commentary presented a more technical view on AI regulation. Some of the assumptions that current AI regulation makes, particularly regarding LLMs and similar models, are fiercely debated within the AI community. Therefore, it is important that in the future, computer scientists and AI researchers take a more active role in policy debates. This is a contribution to this debate.

References

- [1] African Union Development Agency. 2023. *AUDA-NEPAD White Paper: Regulation and Responsible Adoption of AI in Africa Towards Achievement of AU Agenda 2063*. Technical Report.
- [2] Marco Almada and Nicolas Petit. 2023. The EU AI Act: a medley of product safety and fundamental rights? *Robert Schuman Centre for Advanced Studies Research Paper* 2023/59 (2023).
- [3] Marco Almada and Anca Radu. 2023. The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy. *German Law Journal* (2023), 1–18.
- [4] Patrick Altmeyer, Andrew M Demetriou, Antony Bartlett, and Cynthia Liem. 2024. Position Paper: Against Spurious Sparks-Doveling Inflated AI Claims. *arXiv preprint arXiv:2402.03962* (2024).
- [5] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 610–623.
- [6] Anu Bradford. 2020. *The Brussels effect: How the European Union rules the world*. Oxford University Press, USA.
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. *arXiv:2005.14165 [cs.CL]*
- [8] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. Sparks of Artificial General Intelligence: Early experiments with GPT-4. *arXiv:2303.12712 [cs.CL]*
- [9] Miriam C Buiten. 2019. Towards intelligent regulation of artificial intelligence. *European Journal of Risk Regulation* 10, 1 (2019), 41–59.
- [10] W Choi, M van Eck, Hooghiemstra T Cécile van der H, and E Vollebregt. 2022. Legal analysis: European legislative proposal draft AI act and MDR/IVDR. *Hooghiemstra and Partners* (2022).

- [11] Roberto Colom and Paul M Thompson. 2011. Understanding human intelligence by imaging the brain. *The Wiley-Blackwell Handbook of Individual Differences* (2011), 330–352.
- [12] Council of the European Union. 2022. 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 - General approach. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=consil:ST_15698_2022_INIT.
- [13] Maarten den Heijer, Teun van Os van den Abeelen, and Antanina Maslyka. 2019. On the use and misuse of recitals in European union law. *Amsterdam Law School Research Paper 2019-31* (2019).
- [14] Lilian Edwards. 2022. *Regulating AI in Europe: four problems and four solutions*. Technical Report. Ada Lovelace Institute.
- [15] Olivia J Erdélyi and Judy Goldsmith. 2018. Regulating artificial intelligence: Proposal for a global solution. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 95–101.
- [16] Wolfgang Ertel. 2018. *Introduction to artificial intelligence*. Springer.
- [17] European Commission. 2021. 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>.
- [18] European Parliament. 2023. Amendments adopted by the European Parliament on 14 June 2023 on the proposal for a regulation of the European Parliament and of the Council on laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. https://www.europarl.europa.eu/doceo/document/TA-9-2023-0236_EN.html.
- [19] European Parliament. 2024. Amendment 808 adopted by the European Parliament on 13 March 2024 on the proposal for a regulation of the European Parliament and of the Council on laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. https://www.europarl.europa.eu/doceo/document/A-9-2023-0188-AM-808-808_EN.pdf.
- [20] Juliette Faivre. 2023. The AI Act: Towards Global Effects? Available at SSRN 4514993 (2023).
- [21] Steven Feldstein. 2023. Evaluating Europe’s push to enact AI regulations: how will this influence global norms? *Democratization* (2023), 1–18.
- [22] Henry Fraser and José-Miguel Bello y Villarino. 2023. Acceptable risks in Europe’s proposed AI act: reasonableness and other principles for deciding how much risk management is enough. *European Journal of Risk Regulation* (2023), 1–16.
- [23] Patrick Glauner. 2022. An assessment of the ai regulation proposed by the european commission. In *The Future Circle of Healthcare: AI, 3D Printing, Longevity, Ethics, and Uncertainty Mitigation*. Springer, 119–127.
- [24] Oskar Josef Gstrein. 2022. European AI Regulation: Brussels Effect versus Human Dignity? *Zeitschrift für Europarechtliche Studien (ZEuS)* 4 (2022).
- [25] Philipp Hacker. 2020. AI regulation in Europe. Available at SSRN 3556532 (2020).
- [26] Paul Henman. 2020. Improving public services using artificial intelligence: possibilities, pitfalls, governance. *Asia Pacific Journal of Public Administration* 42, 4 (2020), 209–221.
- [27] Marius Hobbhahn and Tamay Besiroglu. 2022. Trends in GPU Price-Performance. <https://epochai.org/blog/trends-in-gpu-price-performance> Accessed: 2024-03-20.
- [28] Wolfgang Hoffmann-Riem. 2020. Artificial intelligence as a challenge for law and regulation. *Regulating artificial intelligence* (2020), 1–29.
- [29] Niklas Kühl, Max Schemmer, Marc Goutier, and Gerhard Satzger. 2022. Artificial intelligence and machine learning. *Electronic Markets* 32, 4 (2022), 2235–2244.
- [30] Lan Li, Tina Lassiter, Joohee Oh, and Min Kyung Lee. 2021. Algorithmic Hiring in Practice: Recruiter and HR Professional’s Perspectives on AI Use in Hiring. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (Virtual Event, USA) (AIES ’21)*. Association for Computing Machinery, New York, NY, USA, 166–176. <https://doi.org/10.1145/3461702.3462531>
- [31] Peter JF Lucas and Linda C Van Der Gaag. 1991. *Principles of expert systems*. Addison Wesley Longman.
- [32] Xinyin Ma, Gongfan Fang, and Xinchao Wang. 2023. LLM-Pruner: On the Structural Pruning of Large Language Models. In *Advances in Neural Information Processing Systems*, A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 21702–21720. https://proceedings.neurips.cc/paper_files/paper/2023/file/44956951349095f74492a5471128a7e0-Paper-Conference.pdf
- [33] Josh Meltzer and Aaron Tielemans. 2022. The European Union AI Act. *Bruselj: Brookings Institution* (2022).
- [34] Sean Musch, Michael Borrelli, and Charles Kerrigan. 2023. The EU AI Act As Global Artificial Intelligence Regulation. Available at SSRN 4549261 (2023).
- [35] Nils J Nilsson. 1982. *Principles of artificial intelligence*. Springer Science & Business Media.
- [36] Regine Paul. 2023. European artificial intelligence “trusted throughout the world”: Risk-based regulation and the fashioning of a competitive common AI market. *Regulation & Governance* (2023).
- [37] Huw Roberts, Emmie Hine, Mariarosaria Taddeo, and Luciano Floridi. 2023. Global AI governance: barriers and pathways forward. (2023).
- [38] Stuart Russell and Peter Norvig. 2020. *Artificial Intelligence: A modern Approach* (4 ed.). Pearson, Harlow, UK.

- [39] Stuart Russell, Karine Perset, and Marko Grobelnik. 2023. Updates to the OECD’s definition of an AI system explained. <https://oecd.ai/en/wonk/ai-system-definition-update> last accessed 2024-03-18.
- [40] Nathalie A Smuha. 2021. From a ‘race to AI’ to a ‘race to AI regulation’: regulatory competition for artificial intelligence. *Law, Innovation and Technology* 13, 1 (2021), 57–84.
- [41] Eva Thelisson and Himanshu Verma. 2024. Conformity assessment under the EU AI act general approach. *AI and Ethics* (2024), 1–9.
- [42] Alan Turing. 1950. Computing machinery and intelligence. *Mind* 59, 236 (1950), 433–460.
- [43] European Union. 2022. *Interinstitutional Style Guide: 2011*. Publications Office of the European Union. <https://doi.org/10.2830/215072>
- [44] Michael Veale and Frederik Zuiderveen Borgesius. 2021. Demystifying the Draft EU Artificial Intelligence Act—Analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International* 22, 4 (2021), 97–112.
- [45] Yoshija Walter. 2023. The rapid competitive economy of machine learning development: a discussion on the social risks and benefits. *AI and Ethics* (2023), 1–14.
- [46] Joseph Weizenbaum. 1966. ELIZA—a computer program for the study of natural language communication between man and machine. *Commun. ACM* 9, 1 (jan 1966), 36–45. <https://doi.org/10.1145/365153.365168>
- [47] Andrew C Wicks, Linnea P Budd, Ryan A Moorthi, Helet Botha, and Jenny Mead. [n. d.]. Automated Hiring at Amazon. ([n. d.]).
- [48] Wikipedia. 2024. List of Nvidia graphics processing units. <http://en.wikipedia.org/w/index.php?title=List%20of%20Nvidia%20graphics%20processing%20units&oldid=1212927172>. last accessed 2024-03-20.
- [49] Bernd W Wirtz, Jan C Weyerer, and Ines Kehl. 2022. Governance of artificial intelligence: A risk and guideline-based integrative framework. *Government Information Quarterly* 39, 4 (2022), 101685.
- [50] Carole-Jean Wu, David Brooks, Kevin Chen, Douglas Chen, Sy Choudhury, Marat Dukhan, Kim Hazelwood, Eldad Isaac, Yangqing Jia, Bill Jia, et al. 2019. Machine learning at facebook: Understanding inference at the edge. In *2019 IEEE international symposium on high performance computer architecture (HPCA)*. IEEE, 331–344.
- [51] Fuzhao Xue, Yao Fu, Wangchunshu Zhou, Zangwei Zheng, and Yang You. 2023. To Repeat or Not To Repeat: Insights from Scaling LLM under Token-Crisis. In *Advances in Neural Information Processing Systems*, A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 59304–59322. https://proceedings.neurips.cc/paper_files/paper/2023/file/b9e472cd579c83e2f6aa3459f46aac28-Paper-Conference.pdf
- [52] Pär J. Ågerfalk. 2020. Artificial intelligence as digital agency. *European Journal of Information Systems* 29, 1 (2020), 1–8. <https://doi.org/10.1080/0960085X.2020.1721947> arXiv:<https://doi.org/10.1080/0960085X.2020.1721947>

A Definitions of AI (Systems)

Table 1. Definitions of AI (Systems) discussed in this article

Definition	Source	Year	Citation
“Many human mental activities such as writing computer programs, doing mathematics, engaging in commonsense reasoning, understanding language, and even driving an automobile are said to demand ‘intelligence.’ Over the past few decades, several computer systems have been built that can perform tasks such as these. [...] We might say that such systems possess some degree of artificial intelligence”	“Principles of artificial intelligence” by Nilsson	1982	[35]
“AI is typically used to refer to (systems developed with) machine learning algorithms”	“Improving public services using artificial intelligence: possibilities, pitfalls, governance.” by Henman	2020	[26]

Table 1. Definitions of AI (Systems) discussed in this article (cont'd)

Definition	Source	Year	Citation
<p>“Artificial intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to process information and function similarly to human beings. In mimicking humans’ actions, AI-based systems are exhibiting traits that are associated with a human mind such as learning and problem-solving. AI is characterised by its ability to rationalize and take actions that have the best probability of accomplishing a specific goal. Machine learning (ML), as a subset of AI, deals with computer programmes that can automatically learn from and adapt to new data without being assisted by humans. On the other hand, deep learning, as a subset of AI, involves techniques that can enable automatic learning by absorbing quantities of unstructured data such as text, images, and videos.”</p>	African Union Development Agency	2023	[1]
<p>“software that is developed with one or more of the techniques and approaches listed in Annex I and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with.”</p> <p>[With Annex I:]</p> <ol style="list-style-type: none"> (1) “Machine learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning; (2) Logic- and knowledge-based approaches, including knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning and expert systems; (3) Statistical approaches, Bayesian estimation, search and optimization methods.” 	AI Act Draft European Commission	2021	[17]
<p>“a system that is designed to operate with elements of autonomy and that, based on machine and/or human-provided data and inputs, infers how to achieve a given set of objectives using machine learning and/or logic- and knowledge based approaches, and produces system-generated outputs such as content (generative AI systems), predictions, recommendations or decisions, influencing the environments with which the AI system interacts”</p>	AI Act Draft Council of the European Union	2022	[12]

Table 1. Definitions of AI (Systems) discussed in this article (cont’d)

Definition	Source	Year	Citation
“artificial intelligence system (AI system) means a machine-based system that is designed to operate with varying levels of autonomy and that can, for explicit or implicit objectives, generate outputs such as predictions, recommendations, or decisions, that influence physical or virtual environments;”	AI Act Draft European Parliament	2023	[18]
“An AI system is a machine-based system that for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.”	Updated OECD Definition	2023	[39]
“AI system’ means a machine-based system designed to operate with varying levels of autonomy, that may exhibit adaptiveness after deployment and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments”	Final Version of the AI Act	2024	[19]

Received 8 January 2024; revised 10 June 2024; accepted 4 September 2024