More Than Meets The AI: The Hidden Costs of a European Software Law

By Mikołaj Barczentewicz and Benjamin Mueller | December 1, 2021

The EU's proposed Artificial Intelligence Act (AIA) would create a risk-based framework for regulating AI, with designated "high-risk" sectors subject to a long list of rules that regulate how firms can design, train, and deploy AI systems. The law's definition of AI is critical: It determines what types of software applications must abide by these obligations. However, the law's definition of AI is so sweeping that it would regulate a broad array of software, creating significant costs and seriously damaging the European Commission's Digital Decade ambitions. A narrower definition of an "AI system" will make the AIA more appropriately targeted and, accordingly, less expensive for the European economy.

The law in its current form will touch upon a much wider section of the EU's economy and society than the European Commission publicly states or likely even envisages.

INTRODUCTION

The stated goal of the EU's proposed Artificial Intelligence Act (AIA) is to regulate artificial intelligence (AI) systems.[1] Defining AI, however, is challenging because it is not a clearly demarcated technology. In this report, we show how the AIA defines AI so broadly as to cover most software. The law in its current form will touch upon a much wider section of the EU's economy and society than the European Commission publicly states or likely even envisages. Arthur C. Clarke stated that any sufficiently advanced technology is indistinguishable from magic: According to the AIA, any sufficiently advanced software is indistinguishable from AI.

THE AIA'S "ALL SOFTWARE IS AI" APPROACH

The reach of a law regulating AI systems hinges on what the law considers to be AI. This creates a problem: There is no commonly accepted legal definition for AI. AI is an umbrella term used to describe a branch of computer science and can include a wide range of software that has little more in common other than its capability of performing tasks that typically require human intelligence.

Consider the wide variety of different types of AI applications:

- Classification: identifying and assigning characteristics or properties, including image processing and computer vision (e.g., automatic friend tagging on photos, automated number plate recognition, estimating the probability of default on a loan, spotting tumors on MRI scans, scanning handwriting, and turning it into text)
- Continuous estimation/regression: predicting the value of an entity in a series (e.g., house price estimates, forecasting demand for electricity)
- Clustering: sorting and classifying a population based on shared characteristics (e.g., identifying market segments, recommending products to buy, word association)
- Skill acquisition: learning to play games (e.g., AlphaGo), autonomous driving, solving tasks and problems, text generation (e.g., GPT-3).[2]

Or consider the difference in the various AI approaches:

- Supervised learning: data labeled by humans and fed into the algorithm
- Unsupervised learning: data clustered/dimensionality reduced without additional input
- Reinforcement learning: give system an objective, don't tell it how to achieve it, provide feedback (the "reward") on whether objective was achieved or not.[3]

Or consider five different fundamental design philosophies that have been driving the discipline of AI since its inception (table 1).

Table 1: Five AI design paradigms[4]

AI Design Paradigm	Examples
Expert Systems / Symbolic AI	Early chatbots (e.g., ELIZA); Neurosymbolic AI
Artificial Neural Networks	Convolutional Neural Networks (e.g., image recognition); Transformer Networks (e.g., generating text); Recurrent Neural Networks (e.g., natural language processing); Generative Adversarial Networks (e.g., generating images)
Bayesian Networks	Causal inference / directed acrylic graphs
Evolutionary Algorithms	Recommender systems
Analogizers	Kernel machines (e.g., support vector machines)

Defining AI in a way that captures both existing and new approaches, in a constantly evolving technical discipline, is a challenge. So how does the AIA go about it? In Article 3(1) the law defines an "AI system" 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.

The list in Annex I includes:

> (a) Machine-learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning;
>
> (b) Logic- and knowledge-based approaches, including knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning and expert systems;
>
> (c) Statistical approaches, Bayesian estimation, search and optimization methods.

In short, under the AIA, a system is designated as AI if it meets two criteria: 1) it is developed using one of the methods listed in Annex I, and 2) it generates output for a given set of human-defined objectives. The second of these criteria is almost meaningless, since this would only seemingly disqualify software that is either not designed by humans or which generates no output. This first criterion is equally expansive. While Clause (a) lists common machine-learning approaches, Clause (b) expands the law's scope to include an enormous range of software. After all, what software isn't ultimately "logic or knowledge-based"? In addition, the inclusion of statistical approaches in Clause (c) means that even software using a basic linear regression model meets the AIA's definitional threshold.

According to the AIA's authors, the definition is meant to be "as technology neutral and futureproof as possible, taking into account the fast technological and market developments related to AI." The implications of the AIA's definition of AI become apparent when we think of the myriad software applications the law will deem "high-risk," where the law will impose the most significant regulatory burdens. The sectors in which the use of AI is considered "high-risk" are:

- Critical infrastructure where the AI system could put human life and health at risk;
- Educational and vocational settings where the AI system could determine access to education or professional training;
- Employment, worker management and self-employment;

- Essential private and public services, including access to financial services such as credit scoring systems;
- Law enforcement;
- Migration, asylum and border control, including verifying the authenticity of travel documents; and
- The administration of justice.

A software tool developed using one of the approaches in Annex I and applied in a "high-risk" sector needs to comply with a detailed list of technical and auditing requirements:

- Creating and maintaining a risk management system for the entire lifecycle of the system;
- Testing the system to identify risks and determine appropriate mitigation measures, and to validate that the system runs consistently for the intended purpose, with tests made against prior metrics and validated against probabilistic thresholds;
- Establishing appropriate data governance controls, including the requirement that all training, validation, and testing datasets be complete, error-free, and representative;
- Buying and maintaining a quality management system;
- Maintaining detailed technical documentation, including around system architecture, algorithmic design, and model specifications;
- Automatic logging of events while the system is running, with the recording conforming to recognized standards;
- Designed with sufficient transparency to allow users to interpret the system's output;
- Designed to maintain human oversight at all times and prevent or minimize risks to health and safety or fundamental rights, including an override or off-switch capability; and
- Conducting a conformity assessment to validate compliance with the above requirements.

To understand the implications of this expansive definition of AI, consider the following five examples of basic software tools that would spark the AIA's "high-risk" requirements.

EXAMPLE 1: SCHOOL DEVELOPS ADMISSION SUPPORT SYSTEM

A school develops a simple logic-based expert system to assist in making decisions related to admissions. It could be as basic as a Microsoft Excel macro checking if a candidate is in the school's catchment area, by linking the candidate's postcode with the content of one column of a spreadsheet, and comparing the output with another column. Under the current definitions in the AIA, this would not only be an "AI system", but also a "high-risk AI system" because it is "intended to be used for the purpose of determining access or assigning natural persons to educational and vocational training institutions." (Annex III). Hence, to use this simple Excel macro, the school would be legally required to fulfil the list of requirements for "high-risk AI systems."

EXAMPLE 2: SMALL BUSINESS DEVELOPS RECRUITMENT WEBSITE

A small business develops a website to allow job applicants to apply to open positions. Some positions legally require certain professional licenses. The website uses a basic logic system to alert applicants who do not have the required professional license that they do not meet the qualifications for the position. Based on the criteria in the AIA, this website would qualify as a "logic- or knowledge-based AI system" and because it is an AI system "intended to be used for recruitment or selection of natural persons, notably for advertising vacancies" it would have to meet the requirements of a "high-risk AI system."

EXAMPLE 3: FINANCE BUSINESS ENHANCES CUSTOMER SERVICE

A bank develops a knowledge-based tool to enable its customer service agents to quickly find information relevant to handling an incoming caller's requests. It has designed this system in response to many customer complaints about their calls routinely shuttling from one agent to another. Previously, case information was siloed between different customer service departments and calls had to be routed from agent to agent to resolve complex queries. The knowledge-based tool gives single agents access to all relevant case files. Although this tool empowers customer service agents to handle calls more quickly and efficiently, because it is a knowledge-based tool and relates to access to financial services, it is considered "high-risk" and must meet the AIA's compliance rules.

EXAMPLE 4: GOVERNMENT SERVICES PORTAL UTILISES A RECOMMENDER SYSTEM

To improve the ability of citizens to navigate a complex government website that collates a large number of public services, a basic chatbot-based recommender tool is utilised to rank possible solutions to a problem faced by the user. The aim is to point citizens into the right direction quickly, lowering the time they spend navigating the website; and to reduce call volumes to the government agency. The tool uses a pre-programmed decision tree, which meets the AIA's definition of AI, and because it is used in the context of government services it falls under the "high-risk" category of the AIA and must comply with the law before being deployed.

EXAMPLE 5: CITY USES ROAD PRICING AND TRAFFIC MANAGEMENT SYSTEM

To combat pollution and congestion, a city authority wants to roll out a dynamic road pricing system that charges road users based on the type of vehicle they drive, the time of day, and the area they are driving in. The city authority is also connecting this system to a city-wide traffic management tool that adjusts traffic light duration at various choke points to reduce engine idling and pollution. The system combines Automatic Number Plate Recognition, stationary traffic sensors, and public transport GPS data into a regression model. Even though this system consists of technology that has been available for two decades and hardly counts as artificial intelligence (e.g., London's road pricing system went live in 2003), it meets the AIA's

definitional threshold and relates to critical infrastructure and must therefore undergo the AIA's conformity assessment procedure before it goes live.

Because the AIA sweeps up tools that go far beyond what most would reasonably consider "AI," and because the law includes a vast range of "high-risk" application areas, its requirements will apply widely to all kinds of software solutions. There are, of course, crucial tradeoffs inherent in any legal definition of AI, but these are not acknowledged by the authors of the AIA. Technological neutrality and future-proofing regulation may be valuable goals. However, attempting to achieve this with a definition far beyond what can meaningfully be called AI leads to unintended consequences that dramatically raise the cost of the law.

UNINTENDED CONSEQUENCES AND HIDDEN COSTS

A broad definition of AI means that the AIA will have a much wider cost footprint than the law's supporters admit. In response to a recent analysis from the Center for Data Innovation, which estimates the AI Act to cost the European economy upwards of €30B over the next five years, the Commission and a number of academics have come out in protest. A key counterargument is that the estimate is too high because the analysis applies the cost of compliance for high-risk AI too widely. Both the Commission and various academics have reiterated, without supporting evidence, that only 10 percent of AI systems will be deemed high-risk.[5] Not only does this argument ignore the vagueness and uncertainty surrounding the high-risk definition, but it also fails to account for the latitude of the definition of AI that we highlight in this report.

The AIA's definitional problems will undermine the ability of Europeans to deploy innovative software systems. Rules and definitions that might seem sensible to lawmakers do not align with the technical realities of software development. The imposition of expensive rules for AI systems will particularly discourage small and new organizations from the regulated market. The Commission has noted as much in its own Impact Assessment support study:

> While these estimates assume that high-risk AI systems only count for 10% of total AI investments, the actual proportion is unknown and will depend on the definition of high-risk AI systems. The private sector will also respond to the new regulation and thus the equilibrium high-risk AI investment will be determined endogenously.[6]

The AIA creates prohibitive burdens for the development and adoption of useful digital tools that should be widely available in Europe's "Digital Decade." If a business wants to build a "high-risk" tool that falls under the AIA's definition of AI, the obligatory conformity assessment means that companies without such systems in place need to build them from scratch, with costs of up to €400,000.[7] Building or buying "high-risk" tools will likely cease to be an option for small organizations. As the Study to Support an Impact Assessment notes:

> Some stakeholders interviewed stated that they might refrain from producing any regulated AI systems to avoid additional costs. [...] SMEs may lack significant funds and thus choose to stay away from the regulated market. [...] The one-off [compliance] spend may deter new entrants from developing high-risk AI systems, but have less impact on existing companies that have already made their investment decision.[8]

Moreover, smaller software developers who today create custom software for organizations may not be able to continue offering their services in "high-risk" sectors. This is because every such custom software product may have to go through all the regulatory requirements independently. Hence, a consequence of the AIA may be that instead of embracing technological change, large pockets of the European economy will continue to persist in using low-tech solutions even for such simple and low-risk tasks as the ones described in the case studies above. The opportunity of not just actual AI, but software in general, to increase productivity, wages, and growth will go to waste in large segments of the economy. The "high-risk" designation will affect about one-third of the European economy, creating additional regulatory burdens for many software tools used in such sectors.[9] This is a major roadblock for the EU's "Digital Decade" ambitions. Just as we enter an age of rapid progress in AI, the AIA will push the value creation brought about by new AI tools out of Europe. The extra development costs in Europe that even innocuous software tools in "high-risk" areas incur incentivizes business to build and develop AI products elsewhere and enter the European market only once they can shoulder the costs of the AIA.

Another response to the Center's cost analysis—that building a quality management system is business as usual, and few companies will need to do so from scratch—emerges from a static assumption that the law only applies to existing companies, many of whom are already covered by existing safety and product laws. The promise of a greenfield technology like AI is to spur the creation of a generation of new businesses that leverage the technology. The idea that the AIA's obligations primarily affect existing businesses reinforces a key argument of the Center: the AIA will lead to enormous opportunity costs in terms of new companies and tools foregone, and opportunities not seized. The endogeneity effects of the AIA will vastly outweigh the static cost assumptions of the European Commission's Study to Support an Impact Assessment, which looks only at costs to existing businesses, and assumes that most of the regulated market will be seized by legacy players. That assumption of market lethargy in and of itself is a striking symbol of the AIA's cost to the European economy.

The debate around the AIA tends to focus on the obligations it creates for high-risk AI systems. However, any company or entity that offers an AI-based service to a European customer is covered by the law, regardless of the risk category in which it falls. Risk categories, of course, may change in the future via delegated acts, and a sector currently classed as "low risk" could eventually become "high-risk." Any prudent organization engaged in what the law deems "artificial intelligence," no matter what risk category it is in, will

need to study the AIA and prepare for the eventuality of the law's risk categories widening in future. The possibility that the AIA will be interpreted and even changed to impose greater regulatory burdens is strengthened by the fact that the AIA is drafted as a harmonizing regulation. It will preempt at least some national laws within the AIA's material scope—i.e., on "the placing on the market, the putting into service and the use of" nearly all software. With such an ambitious scope, the pressure to use the AIA to address any politically salient issues of the day may be great.

The AI Act departs from previous approaches to technology regulation whereby organizations roll out a new technology or service, go to market, and adapt to regulations that are crafted to respond to identifiable risks and problems. The AIA, by contrast, requires companies to abide by a set of rules before they deploy AI-based products, rules that are untested and unsupported by evidence establishing their impact and utility. This *ex ante* regulatory approach, which has not been tried on technologies as ubiquitous and varied as software, expands the law's financial burden and will slow down the digitization of the European economy.

MORE DEFINITIONAL PROBLEMS IN THE AIA

Finally, we want to highlight two more definitional quagmires that further raise the law's cost.

Firstly, the AIA's prohibition of "subliminal systems" is a conceptual minefield:

> Placing on the market, putting into service or using an AI system that deploys subliminal techniques beyond a person's consciousness to materially distort a person's behavior in a manner that causes that person or another person physical or psychological harm.

The terms "subliminal techniques", "beyond a person's consciousness", "materially distort a person's behaviour" and "psychological harm" are new, open-ended phrases that are undefined in EU law. They harken back to the fears around television advertising in the 1960s that polemicist Vance Packard warned about.[10] Indeed, isn't "materially distorting a person's behaviour" what all marketing and advertising seeks to do (in this case, encouraging someone to buy something)? These terms' meanings will be determined by future legislation and case law. If judges subsequently give these concepts a broad interpretation, the AIA could inadvertently end up covering broad swathes of the digital economy including many forms of marketing and advertising as well as user-interface design. Already, activists and some Commission officials are suggesting that this clause could end up prohibiting some forms of online advertising.[11] Plenty of curated content provision on newsfeeds and recommender systems that are a common feature on social media platforms could fall under this definition and thus be banned. Moreover, audiovisual content and games may produce unintentional "subliminal" effects that affect user behavior. The use of unscientific and vague concepts like "subliminal" in a binding legal text is concerning. It is imperative that lawmakers provide clearer, and more precise definitions of these concepts to provide certainty to businesses regarding

what the law will regulate. Similarly, the Act deems "high-risk" all uses of "remote biometric authentication" that take place "without prior knowledge of the user" and "at a distance". All of these terms need to be clarified. Does a system that temporarily detects and categorizes facial or physical characteristics but does not identify an individual (for example, an AI system that measures overall customer flow within a store) count as a "remote biometric identification system"? This showcases how the AIA's reliance on horizontal definitions—"AI systems intended to be used for the 'real-time' and 'post' remote biometric identification of natural persons pose significant risks to the health and safety or fundamental rights of persons"—sweeps up a whole host of innocent, user-friendly digital tools that make our lives easier.

WHAT SHOULD BE DONE?

The second-order costs associated with a complex, single-subject "software law" are not properly addressed by European policymakers. Moreover, it is impossible to properly discern the policy concerns and to design appropriate legislative mitigation measures this early in the development cycle of a new technology—comparable to regulators wanting to put in place a full-flung set of automobile engine design and safety specifications, as well rules for road transportation and driver certification, in 1910.

Pointing out the costs of an enormous one-size-fits-all law is not an argument against any regulation. Instead, the questions we raise are: What are the specific ills the AI law is supposed to remedy or prevent? How can these outcomes be achieved at a lower cost? What provisions and burdens of the law can be reduced in scope, raising the economic benefits AI can generate without increasing risks to consumers?

The underlying problem behind the AIA, and the reason why the law ends up regulating most modern software through a definitional backdoor, highlights an inherent confusion behind the law's purpose. Is the objective to control software engineering and development? That is the law's current implication, whether by accident or design. If the AIA's objective is to control only certain AI software engineering practices that pose unique risks to European citizens, then it should focus on specific, identifiable risks based on a clearly defined risk framework—rather than the fraught approach of vaguely designating specific sectors of the economy as "high-risk" in the context of AI systems, or, as we have argued, software.

Specifically, if the AIA's purpose is to achieve consumer protection against specific AI-incurred risks, lawmakers should delineate those specific risks that existing laws fail to account for and focus on regulating the issues that they believe AI will cause, instead of regulating the technology itself. For example, if the concern is that firms are making unfair hiring decisions, then a better policy is surely to create more oversight and accountability over hiring, regardless of whether firms use AI or not, rather than merely regulating AI systems used for hiring and ignoring the macroeconomic cost of such an approach.

The decision to define AI via an extremely broad list of technical approaches reflects the ambiguities behind the purpose of the law. A definition of AI that includes any product built based on a specific list of software design paradigms is going to lead to a degree of regulatory intervention in the digital economy that will substantially raise the cost of doing business in Europe.

We believe that there are two broad options open to policymakers to reduce the definitional minefield that the AIA creates, which will reduce the law's cost without reducing regulatory protection of consumers.

1. **Limit the definition to "black-box" machine-learning paradigms that generate outputs which cannot be understood or traced by humans.** This approach would entail removing clauses (b) and (c) from the existing definition and scoping clause (a) to "uninterpretable machine learning systems." The AIA would apply only to the class of AI techniques whose models, by their very design, are opaque and inscrutable to human-level explanations. This would incentivize the pursuit of research into interpretable AI methods, though it would penalize AI systems that are reliable and accurate despite being unexplainable.

2. **Replace the current definition with one focusing on a system's capabilities.** If the AIA's purpose is to achieve consumer protection against specific AI-incurred risks, it would make more sense to define the technology via capabilities, not a list of software approaches. For instance, if AI is defined as "a computerized system that can autonomously arrive at decisions which affect the physical welfare of its users," it follows that such systems need to fulfil certain criteria to be admissible. There is precedent for such definitions: In the United States, the National Institute of Standards and Technology defines AI as:

> (1) A branch of computer science devoted to developing data processing systems that performs functions normally associated with human intelligence, such as reasoning, learning, and self-improvement.
>
> (2) The capability of a device to perform functions that are normally associated with human intelligence such as reasoning, learning, and self-improvement.[12]

The AIA, however laudable its aims, casts the net too widely in terms of what it regulates. The law's definition should be narrowed to apply only to systems that can learn by themselves, adapt over time, and make decisions based on that learning similar to human-level cognitive acts (such as driving a car, or making judgments about someone's job application). This requires a philosophical shift, away from the top-down *dirigisme* long favored by EU policymakers, and towards a nimbler output-oriented approach that focuses on regulating outcomes, not on the *ex ante* control of technological design. The latter approach will likely engender economic costs orders of magnitude greater than the former and would cripple the EU's prospects of being a frontrunner in the digital economy. Sadly, this runs directly counter to the

interest of European policymakers, academics, and regulatory rent-seekers whose aim is to use the AIA to subject Europe's digital future to a system of centralized command and control.

REFERENCES

1. 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 (2021) https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206.
2. Bernard Golstein, "A Brief Taxonomy of AI," 10 October 2018, https://www.sharper.ai/taxonomy-ai/.
3. Ibid.
4. Pedro Domingos, "The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World" (London: Penguin Books, 2015).
5. Cristina Lago, "European Commission says €31bn cost estimate for AI rules is 'flawed,'" TechMonitor, 28 July, 2021, https://techmonitor.ai/technology/ai-and-automation/european-commission-says-e31bn-cost-estimate-for-ai-rules-is-flawed; Meeri Haataja and Joanna Bryson, "What Costs Should We Expect from the EU's AI Act?" (working paper, 27 August, 2021 https://osf.io/preprints/socarxiv/8nzb4/.
6. European Commission Directorate-General for Communications Networks, Content and Technology, *Study To Support An Impact Assessment Of Regulatory Requirements For Artificial Intelligence In Europe* (Brussels: European Commission, April 2021), 143, http://dx.doi.org/10.2759/523404.
7. Benjamin Mueller, "How Much Will the Artificial Intelligence Act Cost Europe?" *Center for Data Innovation*, July 2021, 8 https://www2.datainnovation.org/2021-aia-costs.pdf.
8. European Commission, *Study to Support an Impact Assessment*, 155-160.
9. Benjamin Mueller, "How Much Will the Artificial Intelligence Act Cost Europe?" 6.
10. Vance Packard, *The Hidden Persuaders*, (New York: Pocket Books, 1957).
11. Natasha Lomas, "Inside a European Push to Outlaw Creepy Ads," *TechCrunch*, 21 October 2021, https://techcrunch.com/2021/10/21/inside-a-european-push-to-outlaw-creepy-ads/.
12. National Institute of Standards and Technology, "Definition of Artificial Intelligence", October 28, 2019, https://csrc.nist.gov/topics/technologies/artificial-intelligence.

ABOUT THE AUTHORS

Dr. Mikołaj Barczentewicz is a senior lecturer in law and the research director of the Law and Technology Hub at the University of Surrey. Dr. Barczentewicz is also a Research Associate of the University of Oxford Centre for Technology and Global Affairs and a Fellow of the Stanford Law School and University of Vienna Transatlantic Technology Law Forum. His research spans technology law and policy, UK and EU public law, and legal philosophy.

Benjamin Mueller is a senior policy analyst at the Center for Data Innovation, focusing on AI and technology governance. Prior to joining the Center he was chief of staff at a financial technology company in London. Dr. Mueller studied Politics, Philosophy and Economics at the University of Oxford, and completed his PhD in International Relations at the London School of Economics.

ABOUT THE CENTER FOR DATA INNOVATION

The Center for Data Innovation is the leading global think tank studying the intersection of data, technology, and public policy. With staff in Washington, D.C., and Brussels, the Center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as open data, artificial intelligence, and the Internet of Things. The Center is part of the nonprofit, nonpartisan Information Technology and Innovation Foundation (ITIF).

Contact: info@datainnovation.org

datainnovation.org