An Assessment of Key Terms in the EU Artificial Intelligence Act

Roger Clarke

Review Draft of 26 July 2026

Abstract

The Artificial Intelligence Act (AIA) of the European Union (EU) has been portrayed as a world-leading regulatory regime that will protect the public against the technological threats inherent in AI, and encourage adoption of beneficial applications. A great deal of analysis has been published since the law was proposed in 2021. This has continued following its passage in 2024, and the coming into force of the first of its provisions in 2025. This article is a contribution to the evaluation of the AIA not as law, but as the underpinnings of a regulatory regime -- or, more specifically, of several separate regimes for different categories of application. The particular focus of this article is an assessment of the EU's definitions of the statute's key terms. It is concluded that they represent substantial barriers to understanding by AI practitioners. This is argued to undermine the effectiveness of the AIA as a basis for regulatory regimes.

1. Introduction

In recent years, rapid adoption has been evident of several forms of Artificial Intelligence (AI) that appear to embody both significant benefits and significant threats[1],[2],[3]. Reflecting the extent and seriousness of the threats, widespread concern exists about AI's impacts, and there is demand worldwide for regulatory action[4]. In an endeavour to placate widespread public concern about the threats, statutory intervention was developed by the European Commission (EC) in the form of the Artificial Intelligence Act (AIA)[5]. This was enacted in 2024, with its provisions coming into effect in a series of steps.

The AIA is claimed to be the first substantive endeavour to regulate applications of AI technologies, and is progressively coming into effect between March 2024 and August 2028. The EC seeks to encourage development, deployment and use of Artificial Intelligence, in order to extract economic benefits from it. Both extensions and further exemptions have already been granted in respect of multiple features, and it appears reasonable to expect that industry lobbying will continue to have successes. Many writers

1 Tais Fernanda Blauth, Oskar Josef Gstrein and Andrej Zwitter, 'Artificial Intelligence Crime: An Overview of Malicious Use and Abuse of AI' [2022] *IEEE Access* 10, 77110-77122 <https://doi.org/10.1109/access.2022.3191790>.

2 Roger Clarke, 'Why the world wants controls over Artificial Intelligence' (2019) 35(4) *Computer Law & Security Review* 423-433 <https://doi.org/10.1016/j.clsr.2019.04.006>.

3Roger Clarke, 'Principles for the responsible application of Generative AI' (2025) 57 *Computer Law & Security Review* 106131 <https://doi.org/10.1016/j.clsr.2025.106131>.

4Nicole Gillespie, et al. 'Trust, atttudes and use of artificial intelligence: A global study' The University of Melbourne and KPMG, 29 April 2025 <https://doi.org/10.26188/2882291

5European Union, 'Regulation on harmonised rules on artificial intelligence' (EU) 2024/1689, August 2024, at https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

see the AIA as being likely to exhibit the 'Brussels Effect'[6], by influencing the practice of AI well beyond the European Union (EU)[7] -- although some conclude differently[8].

The majority of commentary on the AIA has been positive, or at least hopeful, with relatively few works critically examining the proposal or the political economy of its development, passage and implementation. The purpose of the present paper is to examine key terminology underlying the statute, and assess the degree of correspondence between, on the one hand, the meanings of terms imputed by the drafters and, on the other, the perceptions of the world of AI shared by practitioners of the art, science and practices in that field. This work is motivated by the public policy need for control to be exercised over inherently dangerous and obscure technology.

The AIA has been the subject of a vast amount of analysis since its path towards enactment began in April 2021[9]. It is large and complex, and contains about 90,000 words in 100 pp. of dense text. The Recital alone contains 34,500 words in 180 numbered paragraphs over 31 pp. This has since been augmented by guidance totalling a further 135pp. and 60,000 words, relating to only the first provisions that came into effect[10], with a second of 36pp.[11] published a mere 18 days before the relevant provisions came into effect. The scale and complexity of the text of the AIA inevitably gives rise to a rich diversity of interpretations.

In their endeavours to address a multi-headed and rapidly-changing phenomenon, the AIA's drafters used a number of key terms in novel ways. The phrase 'AI System' occurs 1,080 times, 'risk' qualified in several ways occurs 776 times, 'high-risk' 472 times and 'systemic risk(s)' 80 times. The usages of the key terms, combined with the richness of their contexts of use in the document, present to regulatees as a blizzard of words, and provide researchers and consultants alike with an intellectual banquet on which to feast.

The analysis presented in this article is not legal scholarship, and it has not been performed by a lawyer. The perspective adopted is that of a consultant and researcher in strategic and policy aspects of transformative and disruptive information technologies, and the focus is on law as an instrument of public policy. It reflects insights in the substantial prior literature in the field of technology law, including early analyses[12] ,

---

6Anu Bradford, *Brussels Effect: How the European Union Rules the World* (Oxford University Press, Incorporated, 2020).

7Graham Greenleaf 'The 'Brussels effect' of the EU's 'AI Act' on data privacy outside Europe' Privacy Laws & Business 1 (2021) 3-7, at https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=3898904, Graham Greenleaf 'EU AI Act: The second most important data privacy law' Privacy Laws & Business, June 2024, at https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=4913686

8Martin Ebers, 'Truly Risk-based Regulation of Artificial Intelligence How to Implement the EU's AI Act' [2024] *European Journal of Risk Regulation* 1-20 <https://doi.org/10.1017/err.2024.78>, pp.18-20

9European Commission, 'Proposal for a Regulation on a European approach for Artificial Intelligence' 21 April 2021, athttps://ec.europa.eu/newsroom/dae/document.cfm?doc_id=75788

10European Commission, 'Commission Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689 (AI Act)' Draft of 4 February 2025, at https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-prohibited-artificial-intelligence-ai-practices-defined-ai-act

11European Commission, 'The General-Purpose AI Code of Practice', 10 July 2025, at https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai

12Lilian Edwards 'The EU AI Act: a summary of its significance and scope' Ada Lovelace Institute, April 2022 <https://www.adalovelaceinstitute.org/wp-content/uploads/2022/04/Expert-explainer-The-EU-AI-Act-11-April-2022.pdf>

optimistic contributors[13] and particularly critical works in the business-oriented information systems literature[14] and in law[15].

The article commences with a brief review of AI (section 2), followed by an outline of the AIA, and its declared purposes (section 3). It then examines each of the key terms in the statute (section 4), assessing the understandability by AI practitioners of the definitions and explanations that the EC provides for those terms.

2. Artificial Intelligence

The field of AI was established in 1955. Multiple and diverse AI teechnologies have been conceived. Some have made limited headway, some have achieved modest levels of adoption in specific areas, and a few have entered the mainstream. One early example was logic programming, which was developed from about 1970, and applied during the 1980s to the encoding of statutes and government policies and procedures. Another moderate success originating in the same period has been rule-based expert systems.

Multiple applications of pattern recognition have become mainstream. For example, images are analysed to identify printed characters for the extraction of text (optical character recognition, OCR) and of the identifiers of motor vehicles (automated number/license plate recognition, ANPR/ALPR). There have also been successes in the area of sound, such as music, by means of acoustic fingerprinting techniques.

Two particular AI technologies have been the subject of a great deal of promotion and media coverage during the last decade. Both have matured rapidly. Both are largely based on a particular underlying technique referred to as 'artificial neural networks' (ANN). This shifts away from the longstanding approach adopted to software development, whereby designs reflect an underlying model of the domain of application, and there is clarity about what problems are being addressed, and the what and how of the solutions to those problems. ANN technology instead adopts a purely empirical approach, essentially shovelling large volumes of data into a heterogeneous collection, and using brute-force numerical methods to identify correlations among the data-heap. In most circumstances, the connection between theory and practice has been abandoned in favour of solely pragmatic activities not guided by any body of theory or model of aspects of the real world problem-domain.

ANNs were invented in the 1960s. They depend on access to large volumes of data and very substantial processing power. Both of those capabilities became available from about 2010. The earlier of the two ANN-based AI technologies to emerge is commonly referred to as the AI/ML (machine learning) approach to data analytics. This technique draws inferences from the huge data-heaps that were gathered (in many cases by expropriation)

13 Andrej Zwitter, Oskar Josef Gstrein and Noman Haleem, 'General-purpose AI regulation and the European Union AI Act' [2024] *SSRN Electronic Journal* 1-26 <https://doi.org/10.2139/ssrn.4916400>, Prof Celso Cancela Outeda, 'The EU's AI Act: A Framework for Collaborative Governance' [2024] *Internet of Things* 101291 <https://doi.org/10.1016/j.iot.2024.101291>.

14 Fanny Vainionpaa, et al. 'A Review of Challenges and Critiques of the European Artificial Intelligence Act (AIA)' Proc. Int'l Conf. Infor. Syst., 2023, 14, at https://oulurepo.oulu.fi/bitstream/handle/10024/47651/nbnfioulu-202402061598.pdf, Woersdoerfer M. (2024) 'Mitigating the adverse effects of AI with the European Union's artificial intelligence act: Hype or hope?' Global Business and Organizational Excellence 43,3 (January 2024) 106-126, at https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=4630087

15 Veale M. & Borgesius F.Z. (2021) 'Demystifying the Draft EU Artificial Intelligence Act' Computer Law Review International 22,4 (2021) 97-112, at https://arxiv.org/pdf/2107.03721, Irena Barkane, 'Questioning the EU proposal for an Artificial Intelligence Act: The need for prohibitions and a stricter approach to biometric surveillance' [2022] *Information Polity* 2022, 147-162, <https://doi.org/10.3233/ip-211524>, Wachter S. (2024) 'Limitations and Loopholes in the EU AI Act and AI Liability Directives: What This Means for the European Union, the United States, and Beyond' Yale Journal of Law & Technology 26,3 (2024) 671-718, at https://yjolt.org/sites/default/files/wachter_26yalejltech671.pdf

during the phase of 'big data' euphoria c.2005-2015[16]. This empirical technique has been supplanting longstanding inferencing techniques derived from statistical methods and the 'data mining' approaches of the late twentieth century.

A second ANN-based AI technology burst on the scene in 2022. It is commonly referred to as Generative AI (GenAI), and is associated with the early mover, ChatGPT. This differs from AI/ML in that, rather than drawing inferences from a data-heap, it exploits the data-heap to deliver new data that endeavours to fulfil a request made to it. To extend a popular parody, AI/ML is fed a lot of pictures of cats, to create software that recognises cats; whereas GenAI is fed a lot of pictures of cats, to create software that synthesises new pictures of cats.

Some AI applications are impactful on various interests, including personal freedoms, and have excited public concern. Responses by parliaments and regulators have been very limited to date, however. For example, some jurisdictions have created weak codes of practice in relation to the use of ANPR/ALPR and to the incorporation of facial recognition into closed circuit television (CCTV). The EC's proposals, and the drafting of the AIA, were therefore pioneering work, which, in common parlance, invited arrows in their back. The analysis conducted below in effect fires a volley of arrows.

---

3. The AI Act

The paucity of precedents for the regulation of AI technologies was only one challenge that confronted the AIA's drafters. Their instruction was to establish regulation of AI in the abstract, or as a whole, rather than just particular instances of it. The examples offered above are highly diverse in techniques, functions, applications and impacts. Moreover, a degree of public agitation already existed in relation to recently-emerged AI/ML tools, and grew quickly in relation to GenAI during the course of the drafting project.

The AIA was prepared during 2021-24 by the European Commission (EC, which comprises over 30,000 public service employees), passed in March 2024 by the EU Parliament (which in a sense represents citizens but is dominated by the EC), and adopted in May 2024 by the Council of the EU (which represents the governments of the member states). The AIA is formally an EU Regulation, 2024/1689, and hence is a law that applies to and within all member states, as part of each member state's law. (EU Directives, on the other hand, set goals for member states to implement). The AIA became EU law on 1 August 2024, with various elements taking effect from 2 February 2025, and from 2 August in each of 2025, 2026, 2027 and 2028 (Article 113).

The description in this and the following sections draws on the AIA and the EU's high-level summary[17] and a variety of secondary sources in refereed literatures which are cited where they are relied upon. The remainder of this section outlines the AIA's purposes. Section 4 then discusses definitional aspects of the key terms it uses.

AIA's objective is declared as follows (AIA, Recital (1), emphases added, and reformatted for clarity):

The purpose of this Regulation is

- ***to improve the functioning of the internal market** by laying down a uniform legal framework in particular for the development, the placing on the market, the putting into service and the use of artificial intelligence systems (AI systems) in the Union, in accordance with Union values,*

---

16 Roger Clarke, 'Quasi-Empirical Scenario Analysis and Its Application to Big Data Quality' Proc. 28th Bled eConference, Slovenia, 7-10 June 2015, PrePrint at http://www.rogerclarke.com/EC/BDSA.html

17 European Commission, 'High-level summary of the AI Act' 27 February 2024, at https://artificialintelligenceact.eu/high-level-summary/https://artificialintelligenceact.eu/high-level-summary/

- ***to promote the uptake** of human centric and trustworthy artificial intelligence (AI)*
- *while **ensuring** a high level of **protection** of health, safety, fundamental rights** as enshrined in the Charter of Fundamental Rights of the European Union (the 'Charter'), including democracy, the rule of law and environmental protection,*
- ***to protect against the harmful effects of AI systems** in the Union, and*
- ***to support innovation**.*

The driver for the AIA is clearly the stimulation of economic development by supporting AI innovation and promoting its uptake. Control over negative impacts, implications and risks is expressed firstly as a constraint on the economic objective ("while ensuring") and then as an apparently secondary objective ("to protect") --which is nominally an objective, but is actually already encompassed by the preceding constraint. Hence, although the large majority of the AIA's provisions relate to protections, the ultimate purpose of the AIA is to overcome barriers to the adoption of AI technologies. The strangeness of the approach adopted by the drafters has been noted by others[18]:

> *"The proposal mixes reduction of trade barriers with broad fundamental rights concerns in a structure unfamiliar to many information lawyers ... [which] brings a range of novelties and tensions"*

On a more positive note, many regulatory regimes focus on the achievement of public trust (i.e. by whatever means), whereas the AIA contains few mentions of 'trust', but 19 of 'trustworthiness', implying an endeavour to earn public trust rather than relying largely on 'public education', in the sense of public relations and marketing.

---

4. Key Terms Used in the EU AIA

This section assesses the degree of correspondence between the meanings of key terms as defined and explained in the AIA and the EC's guidance documents, and as they are understood by AI practitioners.

4.1 An AI System

The drafters sought to avoid having to define AI. Given the diversity of forms and interpretations over its first seven decades, that preference is unsurprising: AI means too many things to too many people and hence skirting around its meaning could be seen an appropriate strategy. However, the drafters' choice of the subject of regulation is a compound noun, including 'AI', defined in terms of particular characteristics or features that are perceived to make them a suitable matter for the imposition of regulatory measures.

The key term is defined as follows (Art. 3(1), emphases added):

> *An **"AI system"** is a machine-based system that is designed to operate with varying levels of autonomy and 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*

AI researchers, developers, vendors and adopters need to understand the AIA's technological scope in order to gauge the extent to which their technology has the attributes that the definition of "AI system" refers to. In effect, the AIA challenges AI

---

18 Michael Veale and Frederik Zuiderveen Borgesius, 'Demystifying the Draft EU Artificial Intelligence Act — Analysing the good, the bad, and the unclear elements of the proposed approach' (2021) 22(4) *Computer Law Review International* 97-112 <https://doi.org/10.9785/cri-2021-220402>, p.98.

practitioners to infer what the undefined term 'AI' means when prepended to 'system', and to infer meaning in a manner consistent with whatever courts determine it to mean, in due course, but perhaps many years hence.

Achieving a reliable interpretation of the Art.3(1) definition turns out to be quite difficult. Multiple features are optional, including 'autonomy' ("varying levels"), 'adaptiveness' ("may exhibit"), 'objectives' ("explicit or implicit") and any particular form(s) of 'output' ("such as"). The remaining, definitional attributes appear to be (i) 'machine-based', (ii) '(some) autonomy', (iii) 'objectives' and (iv) 'inference from input to generate output'.

It would seem that 'machine' is not intended in its original sense of (paraphrasing OED IV.6.b) 'a device with multiple parts with defined functions, involving mechanical or electrical power in the performance of defined work', but rather a device utilising electrical phenomena for the handling of data, without any requirement that force be applied to any real-world object, i.e. a computer. The following interpretation is therefore suggested as being equivalent to, and more readily understood and applied:

> *An "**AI system**", **operationally defined**, is a computer-based set of interacting processes that has some level of autonomy and some sense of objectives, and that draws inference(s) from input to generate output*

A first important consideration is that the notion of 'inference' does not depend on the involvement of a high-level intelligence. All that matters is that input-conditions exist that give rise to an output. Multiple generations of software development tool exist[19]. The process whereby the output is generated can be expressed in various ways, including as machine-language or assembler coding (used exclusively c.1940-1960, and still used for specific purposes), in algorithmic or procedural form (the dominant approach to software development c.1960-c.2000, and in ongoing use), as antecedent-consequent / logical rules (as used in logic programming languages and rule-based expert systems since c.1970 and c.1980 respectively), or using a purely empirical approach such as artificial neural networks (in laboratories since 1960 and in practice since c.2010). Each of them infers from input to generate output, i.e. each satisfies definitional attribute (iv).

In addition, any item of software can be interpreted as having at least implicit objectives (iii). Each has a delegation to perform its pre-programmed functions and hence operates with a (perhaps low) level of autonomy (ii), and runs in one or more computing devices and is thereby machine-based (i).

The AIA definition of an AI system is therefore not limited to any particular generation of development tool and requires no particular attributes of coding techniques or outcomes that were not already apparent in (at the latest) the very first administrative system (bakery valuations) in November 1951[20], which first ran some years prior to the time when the OED recognises the terms 'artificial intelligence' (1955) and 'AI' (1963).

The definition of AI system adopted in the AIA can therefore be argued to fail the test that the Commission set for it, viz.: "the definition should be based on key characteristics of AI systems that distinguish it [sic] 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" (AIA 2024, Recital (12), emphasis added). Others have reached similar conclusions, e.g. that[21] "the AI Act applies not only to machine learning, but also to logic- and knowledge-based approaches (recital 12 AI Act).

---

19 Roger Clarke, 'A contingency approach to the application software generations' (1991) 22(3) *ACM SIGMIS Database: the DATABASE for Advances in Information Systems* 23-34 <https://doi.org/10.1145/126743.126749>.

20 Frank Land, 'Remembering LEO' in *IFIP Advances in Information and Communication Technology* (Springer Berlin Heidelberg, 2012) 22-42 <https://doi.org/10.1007/978-3-642-33899-1_2>.

21 Martin Ebers, 'Truly Risk-based Regulation of Artificial Intelligence How to Implement the EU's AI Act' [2024] *European Journal of Risk Regulation* 1-20 <https://doi.org/10.1017/err.2024.78>, p.12.

As a result, even deterministic software systems [may be] subject to the highest requirements".

In order to contrast the AIA's definition against past and present understanding among AI practitioners of what the term AI means, the remainder of this section considers two definitions, reflecting firstly the original conception and secondly contemporary usage.

The original conception of AI can be reasonably depicted by a paraphrase of McCarthy et al.'s 1955 description[22]:

> ***Artificial Intelligence*** *is accurate simulation on a computer of all aspects of learning and [human] intelligence more generally*

AIA's definition of an "AI system" incorporates aspects of computing, autonomy, objectives, inferencing, outputs, decisions, and learning (a related notion to adaptability), but lacks any sense of a simulation of the integral whole implied by the term 'intelligence'.

That original conception has come to be referred to as 'artificial general intelligence' (AGI) or 'strong AI', because it is recognised as 'aspirational' (or, less charitably, motivational but unachievable). The AGI notion is widely regarded as being inconsistent with AI as it is practised in the 21st century. The large majority of work in the field has long since adopted the approach that human features that contribute to intelligence do not reflect an 'aspiration', but rather are 'inspirational'[23]. An exemplar of this is the metaphorical use of '(artificial) neural networks' to refer to the most common base for the machine learning branch of AI.

The following paraphrase of multiple sources[24] has been proposed as an operational definition of AI as practised during the last three decades[25]:

> ***Intelligence is exhibited by an artefact if*** *it:*
>
> *(1) evidences (a) perception, and (b) cognition, of relevant aspects of its environment;*
>
> *(2) has goals; and*
>
> *(3) formulates actions towards the achievement of those goals;*
>
> *(4) but also, for some commentators at least, implements those actions*

The AIA's definition of 'AI system' is far removed from this interpretation. AIA requires only weak forms of perception (1a), of goal-driven formulation of action (2 and 3) and of autonomous action (4), and no element at all of cognition (1b). The AIA notion is therefore not only much broader in terms of software development tools (as discussed earlier in this section), but also much narrower in terms of AI practitioners' sense of AI-ness (as discussed immediately above).

These significant differences between the AIA notion and the conception common among researchers, developers, vendors and adopters of contemporary AI creates a strong likelihood that the provisions of the statute will map poorly to the realities of the relevant artefacts' features. This makes it very likely that the legal requirements will not be comprehensible to the individuals who are intended to comply with them. The fact that there is a great deal of commonality between the AIA's definition and those of the

---

22 John McCarthy, et al. 'A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence' 1955, Reprinted in AI Magazine 27, 4 (2006), at https://www.aaai.org/ojs/index.php/aimagazine/article/viewFile/1904/1802, p.12.

23 Margaret Boden, *A. I.: Its Nature and Future* (Oxford University Press, 2016), Antonio Lieto and Daniele Radicioni, 'From human to artificial cognition (and back): New perspectives on cognitively inspired AI systems' (2015) 33 *Cognitive Systems Research* 145 <https://doi.org/10.1016/j.cogsys.2014.11.001>.

24 JS Albus, 'Outline for a theory of intelligence' (1991) 21(3) *IEEE Transactions on Systems, Man, and Cybernetics* 21,3 May-June 1991 473-509, <https://doi.org/10.1109/21.97471>, Peter Norvig, *Artificial Intelligence: A Modern Approach* (Pearson, 2020), John McCarthy 'What is artificial intelligence?' Department of Computer Science, Stanford University, 2007, at http://www-formal.stanford.edu/jmc/whatisai/node1.html.

25 Roger Clarke, 'The Re-Conception of AI: Beyond Artificial, and Beyond Intelligence' [2023] 4(1) *IEEE Transactions on Technology and Society* 24-33 <https://doi.org/10.1109/tts.2023.3234051>.

Organisation for Economic Cooperation and Development[26] and the Council of Europe[27] suggests that this incomprehensibility problem may be becoming entrenched.

4.2 General-Purpose AI (GPAI) System

Generative AI (GenAI) burst into public prominence during 2022-23, during the middle part of the gestation and negotiation period for AIA, 2021-24. GenAI utilises ANNs to respond to requests to produce what appears to be new data. The EC felt the need to react swiftly to the very rapid adoption of GenAI in 2022-24, and, despite the fluidity of an immature technology, sought to ensure that a particular aspect of GenAI was addressed by the Act.

The AIA defines one key term as (Art.3(66), emphasis added):

> *A "**general-purpose AI (GPAI) system**" means an AI system which is based on a general purpose AI model [GPAI model], that has the capability to serve a variety of purposes, both for direct use as well as for integration in other AI systems*

The AIA defines a further term as (Art.3(63), emphases added):

> *A "**GPAI model**" means an AI model, including when trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable to competently perform a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications. This does not cover AI models that are used before release on the market for research, development and prototyping activities*

A further term, **'GPAI model with systemic risk'** is discussed in section 3.4 below.

The very apparent desire for a 'black-box' / 'technologically neutral' definition has resulted in highly non-specific, even vacuous, expression, which invites an enormous range of interpretations. GenAI artefacts are complex, comprising deft combinations of four elements, the first three of which are forms of AI[28]:

- Natural Language Understanding (NLU) capabilities (to interpret the request);
- the central element of a Large Language Model (LLM)[29] (which uses a vast index, that is expressed in a very simplistic grammar, to extract an abstract model of a response to the request); and
- Natural Language Generation (NLG) functionality (to express the response);

all accessed by means of:

- a 'chatbot' -- a human-computer interface in which a computer interacts with a human in a manner that simulates human communications behaviour.

---

26 Organization for Economic Cooperation and Development, 'OECD Explanatory Memorandum on the Updated OECD Definition of an AI System' OECD Artificial Intelligence Paper No.8, March 2024, at https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/03/explanatory-memorandum-on-the-updated-oecd-definition-of-an-ai-system_3c815e51/623da898-en.pdf, p.4.

27 Council of Europe, 'Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law' , 5 September 2024, at https://rm.coe.int/1680afae3c, p.3.

28 Roger Clarke, 'Principles for the responsible application of Generative AI' (2025) 57 *Computer Law & Security Review* 106131 <https://doi.org/10.1016/j.clsr.2025.106131>.

29 Nikitas Karanikolas et al, 'Large Language Models versus Natural Language Understanding and Generation' in *PCI 2023: 27th Pan-Hellenic Conference on Progress in Computing and Informatics* (ACM, 2023) 278-290 <https://doi.org/10.1145/3635059.3635104>.

The AIA was enacted very early in the maturation process of LLMs, of their combination with other technologies, and of their applications. To address the challenges, the AIA envisages clarifications arising from consultative and advisory bodies. The EU's initial guidance and code have attracted mixed responses from the large corporations that have invested heavily in LLMs. It has been evident from successive EC announcements that industry lobbying has achieved further successes, including material easing of the AIA's provisions in a further Draft Regulation[30]. More may follow.

There is considerable doubt as to whether the notion of 'GPAI' can provide the basis for an efficacious regulatory regime for GenAI, despite it being the particular form of AI that is currently engendering the most public concern.

4.3 Risk

AIA encompasses five categories of "AI systems", four of which are distinguished on the basis of what it refers to as their level of "risk". It specifies regulatory regimes in relation to four of the five categories, exempting from regulation "Minimal risk AI systems". The AIA approach to risk assessment has been the subject of criticism[31].

In the contexts of security, risk assessment and risk management, the term 'risk' is used in a wide variety of ways, referring variously to a general threat, a specific threat, an incident, harm, the likelihood of harm arising from an incident, or the residual likelihood of harm arising taking into account existing safeguards. The EU's cybersecurity agency, ENISA, applies the term 'risk' in a particularly narrow manner, as "a circumstance or event having a potential adverse effect". This is a combination of a generic threat, or circumstance (e.g. lightning) and an instance of a threat, or event (a specific bolt of lightning).

Most coherently, risk refers to[32] (emphases added):

> *the likelihood that a source of hazard will turn into actual harm*
>
> *the perceived likelihood of occurrence of harm arising to an asset as a result of a threatening event impinging on a vulnerability*

The US NIST Glossary entry recognises the blurring of the 'likelihood' criterion through some unclear merger with the notion of harm or 'adverse impacts'[33] (emphases added):

> *A measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence*

The EU Directive on Cybersecurity adopts a related but different approach to US NIST, defining risk as follows[34] (emphases added):

---

30 European Commission, 'Proposal for a Regulation of the European Parliament and of the Council amending Regulations (EU) 2024/1689 and (EU) 2018/1139 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI)' Draft Regulation 2025/0359, 19 November 2025, at https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52025PC0836

31 Claudio Novelli et al, 'AI Risk Assessment: A Scenario-Based, Proportional Methodology for the AI Act' (2024) 3(1) *Digital Society* 1-29 <https://doi.org/10.1007/s44206-024-00095-1>.

32 Martin Ebers, 'Truly Risk-based Regulation of Artificial Intelligence How to Implement the EU's AI Act' [2024] *European Journal of Risk Regulation* 1-20 <https://doi.org/10.1017/err.2024.78>, p.3, Roger Clarke, 'The prospects of easier security for small organisations and consumers' (2015) 31(4) *Computer Law & Security Review* 538-552 <https://doi.org/10.1016/j.clsr.2015.05.004>.

33 National Institute of Standards and Technology, 'NIST Glossary' Computer Security Resource Center, accessed 25 March 2025, at https://csrc.nist.gov/glossary/term/risk

34 European Parliament, 'Directive 2022/2555 on Cybersecurity' 14 December 2022, at https://eur-lex.europa.eu/eli/dir/2022/2555, Art. 6(9).

> *the potential for loss or disruption caused by an incident and is to be expressed as a combination of the magnitude of such loss or disruption and the likelihood of occurrence of the incident*

The AIA adopts a variant of the EU Directive's approach (Art.3(2), emphases added):

> *"**Risk**" means the combination of the probability of an occurrence of harm and the severity of that harm*

This variant leaves open the nature of the combination and the weightings of the two factors. It appears likely that the AIA notion of Risk will require a succession of judicial pronouncements before its meaning and application achieve adequate clarity.

4.4 Systemic Risk

The AIA identifies five categories of AI System distinguished on the basis of the level of risk they entail. The definition of one form of the fifth category, GPAI models, uses the novel term "systemic risk". The AIA definition is as follows (Art.3(65), emphases added):

> ***Systemic risk** means a risk that is specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the Union market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain*

A further term is then defined, as follows (Art.51-1, emphases added):

> ***A GPAI model with systemic risk** is a GPAI model that either:*
>
> *(a) ... has high impact capabilities evaluated on the basis of appropriate technical tools and methodologies, including indicators and benchmarks; or*
>
> *(b) [is] based on a decision of the Commission, ex officio or following a qualified alert from the scientific panel, it has capabilities or an impact equivalent to those set out in point (a) having regard to the criteria set out in Annex XIII.*

Despite the many warnings that exist about high degrees of technological specificity in legislation, the EC seeks to operationalise **a key expression** in terms of the computational power used, "measured in floating-point operations" (FLOPs) (Art.51-2, emphasis added), but with the threshold measure adaptable by regulatory instrument (Art.51-3):

> *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 floating point operations is greater than 10^25*

Use of this measure appears to assume that greater investment in computation necessarily gives rise to more severe harm. This raises the question as to whether a less remote proxy measure for 'severity of negative impact' might be more appropriate.

The lengthy Recitals section provides further explanation of "systemic risks" (in this case plural, rather than singular or generic). They "include, but are not limited to, any actual or reasonably foreseeable negative effects in relation to major accidents, disruptions of critical sectors and serious consequences to public health and safety; any actual or reasonably foreseeable negative effects on democratic processes, public and economic security; the dissemination of illegal, false, or discriminatory content" (Recital 110), and "a general-purpose AI model should be considered to present systemic risks if it has high-impact capabilities, evaluated on the basis of appropriate technical tools and methodologies, or significant impact on the internal market due to its reach. High-impact capabilities in

general-purpose AI models means capabilities that match or exceed the capabilities recorded in the most advanced general-purpose AI models" (Recital 111, emphasis added). The circular nature of the final words of Recital 111 (effectively limiting it to the 'mostest advanced GPAI model(s)') creates challenges to understanding. Endeavours have been made to disentangle the semantic complexities of 'GPAI with systemic risks'[35].

A Code and Guidelines emerged in July 2025[36], but they are heavily process-oriented and lack conceptual clarity. One of the few instances of articulation of a substantive nature is Appendix 1.4 of the Code (p.37), which identifies four categories of "systemic risks: chemical, biological, radiological, and nuclear (CBRN) attacks or accidents ... ; loss of [human] control ... ; large-scale sophisticated cyber-attacks ... ; and harmful manipulation [of] large populations or high-stakes decision-makers through persuasion, deception, or personalised targeting ... ".

The fog surrounding the AIA's definitions of 'risk', as described in the preceding section, is greatly deepened in the case of GPAI models, where the novel concept of 'systemic risk' is created, but not sufficiently clearly articulated to communicate to AI practitioners.

4.5 Categories of Regulatee

The AIA applies to "AI systems" that fit within criteria relating to the risk of harm arising, and that are unable to invoke an escape clause. The high-level summary[37] says that the AIA places "the majority of obligations" on "providers (developers)", but with some also applicable to "operators". The categories of the entities on which the AIA imposes obligations are defined as follows (Arts.3(3,4,8)):

> A "**provider**" is *a natural or legal person, public authority, agency or other body that develops an AI system or a general-purpose AI model or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge*
>
> A "**deployer**" is *a natural or legal person, public authority, agency or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity*
>
> An "**operator**" is *a provider, product manufacturer, deployer, authorised representative, importer or distributor*

In the software and software-as-a-service sectors, a 'developer' performs functions early in a supply-chain, and a wholesaler, a distributor, a re-seller, or indeed the developer themselves adds value further along that chain and hence is a 'provider' in AIA terms. A 'user organisation' may install and operate the software for their own purposes or may depend on a 'service-provider' to do so. Both a 'user organisation' and a 'service-provider' are presumably 'deployers' under the AIA.

Attempts to clarify meanings do not necessarily achieve their aim, e.g. "Those that intend to place on the market [which, in more common business dialect, appears to encompass

---

35 Lee Bygrave and Ronald Schmidt (2025) 'Regulating Non-High-Risk AI Systems under the EU's Artificial Intelligence Act, with Special Focus on the Role of Soft Law' University of Oslo Faculty of Law Legal Studies Research Paper Series No. 2024-10, 29 January 2025, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4997886, pp.8-10.

36 European Commission, 'The General-Purpose AI Code of Practice', 10 July 2025, at https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai, European Commission, 'Guidelines on the scope of the obligations for general-purpose AI models established by Regulation (EU) 2024/1689 (AI Act)' Draft of 18 July 2025, at https://digital-strategy.ec.europa.eu/en/library/guidelines-scope-obligations-providers-general-purpose-ai-models-under-ai-act

37 European Commission, 'High-level summary of the AI Act' 27 February 2024, at https://artificialintelligenceact.eu/high-level-summary/https://artificialintelligenceact.eu/high-level-summary/

developers, re-sellers and service-providers] or put into service [which appears to encompass service-providers but probably also "users (deployers)"]". The subsidiary terms "authorised representative", "importer" and "distributor" (but not "product manufacturer") are defined in conventional manner in Art.3(5)-(7).

The Provider's "**intended purpose**" is defined. It is distinguished from "**reasonably foreseeable misuse**", which is defined as use "not in accordance with its intended purpose" (Art.3(12)-(13)). Thereafter, the term "intended purpose" occurs many times throughout the Articles, whereas 'misuse' appears only in Arts.9, 10, 13 and 14.

The occurrences in the summary of "user" appear to encompass firstly organisations, and secondly individuals applying the AI system in a self-employed professional capacity, i.e. as an independent contractor, sole trader or business partner. The term "end-user" (which appears only twice in each of the AIA and the EC's high-level summary) appears to apply to employees and contractors acting on behalf of an organisation, and individuals acting in a personal capacity. Given that there are few occurrences in the Articles of "user-organisation/deployer" and "user/end-user", it appears that the EC intends to impose no obligations on such entities. This was subsequently confirmed[38].

There is ample scope for uncertainty to arise about the scope of these definitions, and whether all relevant categories of actor are encompassed. After a "provider" has "developed" a product and "placed it on the market", the product may be transacted through chains of organisations before being "used" by a "deployer" for an "intended purpose" and/or a "misuse", or solely for service-provision to other deployers. Clarity is needed about the obligations of, respectively, each organisation in supply chains or networks, and those "deployers" who make the "AI system" available only to "users" within their own organisation and/or "end-users" within or beyond its boundaries.

---

5. Conclusions

The purpose of the work reported in this article was declared as being to examine key terminology underlying the statute, and assess the degree of correspondence between, on the one hand, the meanings of terms imputed by the drafters and, on the other, the perceptions of the world of AI shared by practitioners of the art, science and practices in that field.

A key issue identified in this analysis is the need for individuals and organisations subject to it to be able to discover what provisions of the AIA apply to their work and to their products and services, to understand those provisions, and to relate them to the development and deployment of their products and services. The key terms used in the AIA -- 'AI System', 'GPAI', 'risk' and 'systemic risk' -- are subject to very different interpretations by the EC from the understanding practitioners have of them. The resulting confusion is all the deeper, given that some of the EC's explanations verge on incomprehensibility. The circumstances and nature of such obligations as the AIA imposes, and their allocation to actors in the value-chain, are also highly unclear. It accordingly appears unlikely that the AIA will motivate adjustments of behaviour by developers and deployers and deliver public protections.

A far more appropriate solution may be found by shifting the focus away from both AI in general and GPIA in particular. The AIA is presaged on the assumption that the appropriate subject of the regulatory regime is a technology. Many have expressed serious doubts about the appropriateness of putting particular technologies at the heart of

---

38 European Commission, 'Commission Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689 (AI Act)' Draft of 4 February 2025, at https://digital-strategy.ec.europa.eu/en/library/commission-publishes-guidelines-prohibited-artificial-intelligence-ai-practices-defined-ai-act, p.6, para. (18).

regulatory regimes[39]. Other categories of subjects of regulatory regimes include an industry-sector or -segment, a practice, an impact, an outcome and a profession. The AIA's definition of AI System, at Art. 3(1), when stripped of unhelpful verbiage, has its focus on "a machine-based system that ... infers ... how to generate ... decisions that can influence physical or virtual environments". Others[40] gave argued that the focus should be both narrowed and broadened to "[any] system that generate[s] decisions that can influence physical or virtual environments", or in those authors' terms, all systems whose functions include automated decision-making.

---

39 Lyria Bennett Moses, 'Regulating in the Face of Socio-Technical Change', in Roger Brownsword, Karen Yeung and Eloise Scotford, *Oxford Handbook of Law, Regulation and Technology* (Oxford University Press, 2017).

40 José-Miguel Bello y Villarino et al., 'Are We Regulating the Right Digital Systems? Testing Emerging Artificial Intelligence Frameworks against Real-World Public Sector Systems' (2025) 48(4) *University of New South Wales Law Journal* 1165-1195 <https://doi.org/10.53637/vrjl5367>.