The EU AI Act Definition of AI System: Interpretive Challenges in Real-World Application and the "Navigating the AI System Definition under the AI Act" (NAIS) Framework

María Villalobos-Quesada[1,2,*]

[1]National eHealth Living Lab, Public Health and Primary Care Department, Leiden University Medical Centre, Albinusdreef 2, 2333 ZA Leiden, the Netherlands

[2]Bioethics and Law Observatory, Faculty of Law, University of Barcelona, Avenida Diagonal, 684 08034 Barcelona, Spain

Abstract

This article examines the AI system definition under Article 3(1) of the European Union Artificial Intelligence Act (AI Act) and the interpretive challenges arising from its practical application. It develops the Navigating the AI System Definition under the AI Act (NAIS) Framework, which operationalizes the seven constituents of the AI system definition to support legal and technical assessment of whether a system falls within the scope of the AI Act. The framework is applied to SPIDeRR, a logistic regression system developed to estimate the probability of rheumatic diseases. The analysis shows that several constituents have limited discriminatory capacity, while the assessment of autonomy and inference remains decisive despite the lack of clear qualification criteria. It also highlights the absence of guidance on the cumulative interpretation of the seven constituents. The NAIS Framework provides a structured and transparent approach to support case-by-case assessments by policymakers, regulators, researchers, developers, providers, and deployers.

Keywords

Artificial Intelligence Act, European Union, AI system, logistic regression, AI techniques, machine learning

1. Introduction

In August 2024, the European Union (EU) Artificial Intelligence Act (AI Act) entered into force. The AI Act constitutes the first comprehensive horizontal regulation specifically addressing artificial intelligence (AI) in the Union. The Regulation aims both to stimulate the development and uptake of AI within the EU internal market, and to promote human-centric and trustworthy AI (Recital 1 and Article 1(1), AI Act). According to the High-Level Expert Group on Artificial Intelligence (AI HLEG), trustworthy AI should be lawful, ethical, and robust [1]. In this context, lawfulness requires compliance with the applicable legal framework, including the AI Act.

The AI Act is mainly structured as a product safety regulation, focusing on risks associated with AI systems placed on the market, put into service, and used in the Union (Article 1(2)(a), AI Act). The Regulation was considered necessary to address risks specific to AI systems that were not fully addressed under existing EU legal frameworks [2], including the General Product Safety Regulation (GPSR) and sector-specific legislation such as the Medical Device Regulation (MDR). Since the AI system definition determines the applicability of the AI Act,[1] determining whether a model or a product fulfills that definition is key.

TRUST-AI: The Second European Workshop on Trustworthy AI. Organized as part of the International Joint Conference on Artificial Intelligence - IJCAI/ECAI 2026. August 2026, Bremen, Germany.

*Corresponding author.
m.j.villalobos_quesada@lumc.nl (M. Villalobos-Quesada)
https://nell.eu/member/show/mar-a-villalobos-2 (M. Villalobos-Quesada)
0000-0003-4930-1982 (M. Villalobos-Quesada)

[1]The AI Act also regulates "general purpose AI models", which constitute a distinct regulatory object. General purpose AI models, however, fall outside the scope of this paper.

This article focuses on the interpretation and application of the AI system definition contained in Article 3(1). This is a field in rapid development, currently with only a few theoretical reports published, some of them not peer reviewed [2, 3]. This article illustrates the challenges of applying that definition in practice, going beyond purely theoretical or conceptual analyzes. To this end, this article employs a legal methodology. The seven constituents of the definition established by the European Commission (2025) are analyzed and operationalized into the NAIS Framework. The framework is then applied to a real use case, the SPIDeRR system, which implements a logistic regression model [4, 5]. Although logistic regression models are often considered machine learning techniques [6, 7], some authors argue that they are not "advanced" machine learning or AI techniques [8], and their qualification as AI systems in the context of the AI Act has been challenged by different interest groups [9, 10]. Systems based on logistic regressions are, therefore, interesting for assessing the applicability of the AI system definition of the AI Act because of their borderline characteristics.

This article is structured as follows. First, the SPIDeRR model is presented, highlighting the relevant technical characteristics of the system. Second, the AI system definition contained in Article 3(1) of the EU AI Act is examined using the seven constituents proposed by the European Commission [11]. Third, the constituents are translated into analytical questions to form the Navigating the AI System Definition under the AI Act (NAIS) Framework. Fourth, the NAIS Framework is applied to the SPIDeRR system in order to assess the fulfillment of each constituent. Fifth, the constituents are discussed individually and collectively, highlighting how interpretive challenges affect the assessment of whether a system falls within the AI system definition and, consequently, the applicability of the AI Act. Finally, the main conclusions are presented.

2. The use case: SPIDeRR, a prediction system based on logistic regression

The SPIDeRR system implements a symptom-based logistic regression prediction model, developed to estimate the probability of predefined outcomes, namely the confirmation or *de novo* diagnosis of a set of rheumatic diseases in individuals with musculoskeletal complaints [12, 4, 5]. The model was trained on data obtained through an online questionnaire that collects self-reported symptoms, basic demographic data, and previous diagnoses. The variables included in the model were selected by the researchers, based on clinical relevance, which can be described as acquisition of human-encoded knowledge [13]. The resulting outputs can be presented as probabilistic risk estimates and risk categories, which may support healthcare-related decision-making. The system has not been deployed yet.

During development, the model analyzed the training data to estimate the coefficients for each predictor through parameter optimization. These coefficients reflect the strength and direction of the statistical association between each predictor and the outcome (confirmation or *de novo diagnosis*). For example, symptoms such as joint pain, morning stiffness, or fatigue were assigned coefficients reflecting the strength and direction of their statistical association with the outcomes. Once training is completed, the deployed model remains fixed with respect to its predictor variables, coefficients, and computational structure. During deployment, the system will not learn from new data; for example, it will not derive new coefficients, or modify existing coefficients based on user interactions. Instead, it will execute a predefined mathematical function in which the variables entered by the user will be combined with previously estimated coefficients to calculate a probabilistic output.

Logistic regression is a statistical modeling technique commonly used in biostatistics, epidemiology, and medical applications [4, 6]. Unlike complex machine learning models, the deployed logistic regression model is inherently transparent and interpretable. The predictor variables, estimated coefficients, and computational structure are explicit and can be inspected, allowing the contribution of each variable to the final probability estimate to be understood.

The SPIDeRR system will be deployed under direct human initiation and control. A patient or physician may choose to complete the questionnaire and request a symptom summary or a probabilistic disease estimate. Only after such explicit user action (which can be interpreted as providing consent for processing the data and receiving the output) does the system process the data entered by the user, and generate an output. The only data source are the answers to the online questionnaire. The system does not autonomously collect data from additional sources, initiate independent processing, monitor users, interact autonomously with other systems, or execute actions based on its outputs. In the current foreseeable deployment scenario, physicians using the system will act within an existing doctor-patient relationship, which is governed by national healthcare legislation (e.g., the WGBO in the Netherlands). Patients will be responsible for deciding whether to use the tool and whether to seek medical care if necessary, based on the generated outputs.

3. The AI system definition of the AI Act

> "The notion of 'AI system' in this Regulation should be clearly defined and should be closely aligned with the work of international organisations working on AI to ensure legal certainty, facilitate international convergence and wide acceptance, while providing the flexibility to accommodate the rapid technological developments in this field. Moreover, 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." (Recital 12, AI Act)

The AI Act defines "AI system" in Article 3(1) as follows: "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".

In Article 96(1)(f), the AI Act requires the Commission to develop guidelines on the practical implementation of, *inter alia*, the AI system definition (see also Recital 12, AI Act). The Commission released the Guidelines addressing the AI system definition, hereinafter referred to as the Guidelines, or Guidelines of the European Commission (GEC), in February 2025 [11]. Additionally, some Member States have published national guidelines, including the Netherlands [14]. It is relevant to clarify that guidelines, including the GEC, are not legally binding.

In the Guidelines, the Commission identifies seven main elements of the definition: "(1) a **machine-based system**; (2) that is designed to operate with **varying levels of autonomy**; (3) that may exhibit **adaptiveness after deployment**; (4) and that, for **explicit or implicit objectives**; (5) **infers**, from the input it receives, how to generate **outputs** (6) such as **predictions, content, recommendations, or decisions** (7) that can **influence physical or virtual environments**." Throughout this article, the seven elements identified by the Commission are referred to as the definition's "constituents". This terminology is adopted solely for ease of reference. The constituents differ in their capacity to delimit the scope of the Regulation, as will be discussed in detail in this section.

3.1. Constituents with limited exclusionary effect

The following constituents have limited value in delimiting the applicability of the definition because they are formulated as a basic characteristic of computational systems (constituent 1), without specifying applicable thresholds (constituent 2), in optional terms (constituent 3), in disjunctive terms such that either of two mutually exclusive categories satisfies the condition (constituent 4), through a non-exhaustive list of examples (constituent 6), or in terms of potential rather than actual effects (constituent 7).

Constituent (1) requires a **"machine-based system"**, which refers to AI systems being developed and run on machines (Recital 12, AI Act; paragraph 11, GEC). In other words, AI systems must operate on a machine and be computationally driven (paragraph 11, GEC). Although this constituent establishes a necessary condition, it reflects a basic and general characteristic of software systems, and provides little basis for distinguishing AI systems from other computational systems. The Commission interprets the notion of "machine-based system" broadly, extending it beyond conventional hardware and software to include quantum computing, biological or organic systems with computational capacity (paragraph 13, GEC).

Constituent (2) refers to systems designed to operate with **"varying levels of autonomy"**. Autonomy is understood in the AI Act as having "some degree of independence of actions from human involvement and of capabilities to operate without human intervention" (Recital 12, AI Act). However, the AI Act does not clarify what degree of autonomy is required to qualify as an AI system. If a meaningful level of autonomy is relevant for establishing the scope of the AI system definition, that remains notoriously unclear. Additionally, it remains unclear whether autonomy should be assessed with respect to the system as a whole or to the individual operations it performs. Systems perform different operations throughout their lifecycle, such as data acquisition, parameter estimation, optimization, inference, and output generation, and the degree of human involvement may differ across these operations (Appendix A). Recital 12 attempts to exclude "simpler traditional software systems or programming approaches", such as those systems "designed to operate solely with full manual human involvement and intervention". Does this mean that any degree of automatic processing is sufficient to satisfy the autonomy requirement? If so, do computational systems not, by their nature, perform at least some operations automatically? It becomes then difficult to identify which software systems would fail to satisfy this constituent.

The degree of autonomy cannot be determined solely by examining the technical characteristics of a system. Rather, it also depends on how tasks, control, and decision-making are allocated between the system and the humans interacting with it. After all, the autonomy of an AI system is not an intrinsic property of the technology itself, but the result of design choices through which humans delegate specific tasks and decision-making functions to the system. The Commission recognizes that autonomy is intrinsically linked to human involvement, human intervention, and human-machine interaction (paragraph 16, GEC). However, it remains unclear how this relational understanding of autonomy should be incorporated into the interpretation of Article 3(1), or how it should be operationalized when assessing whether a system is designed to operate with varying levels of autonomy.

Constituent (3) states that an AI system **"may exhibit adaptiveness after deployment"**, indicating that adaptiveness may, but need not, be present after deployment. Adaptiveness can be understood as "self-learning capabilities, allowing the system to change while in use" (Recital 12, AI Act). It is important because adaptiveness can change how the system works and/or the results and outputs it produces. Because of the way the constituent was constructed, it is a non-decisive condition for determining whether a system qualifies as an AI system, as highlighted by the Commission (paragraph 23, GEC).

Constituent (4) refers to AI systems having **"explicit or implicit objectives"**, offering a limited exclusionary effect. According to the Commission, AI system objectives are the internal goals of the tasks performed by the AI system (paragraph 25, GEC). Explicit objectives are those expressly coded or formally predefined. Implicit objectives arise instead from the process of training or from interacting with the environment (paragraph 24, GEC). Because the constituent is formulated in disjunctive terms, the existence of either category is sufficient to satisfy it.

It is relevant to make explicit that the AI Act differentiates between the AI system objectives and the intended purpose (Recital 12, AI Act). The intended purpose takes into account the context and conditions of use, in other words, it is externally oriented (Article 3(12), AI Act).

Constituent (6) intends to provide information about the type of **outputs** that AI systems can generate. It is framed as a list of possible and non-exhaustive output examples, i.e., **"predictions, content, recommendations, or decisions"**, thereby providing no substantial basis for delimiting the scope of the definition.

Constituent (7) establishes that the outputs of AI systems **"can influence physical or virtual environments"**. Notably, constituent (7) links influence exclusively to the outputs generated by AI systems. While the reference to both physical and virtual environments may be helpful in ensuring that the scope of the AI Act is not limited to systems that exert an effect on the physical world, the interpretive difficulty lies mainly in the notions of "influence" and "can influence". Neither the AI Act nor the GEC establish any thresholds or criteria for determining what degree of influence by outputs (or their use) is required for this constituent to be fulfilled. Nor does the AI Act provide any clarification regarding the causal connection between output and its influence on a physical or virtual environment. Finally, the provision appears to require merely the potential or capability to influence an environment, rather than actual or realized influence. As a result, this constituent has a very limited delimiting value.

3.2. Inference

This leaves one remaining constituent. Constituent (5) states that an AI system "infers, from the input it receives". Among the constituents, the inference requirement appears to perform one of the strongest delimiting functions within the definition, particularly because Recital 12 distinguishes inferential capability from "basic data processing": "A key characteristic of AI systems is their capability to infer. This capability to infer refers to the process of obtaining the outputs... and to a capability of AI systems to derive models or algorithms, or both, from inputs or data." Inference then may take place during the development of the AI system (when models or algorithms are derived), or when the system is used (when outputs are obtained).

Recital 12 of the AI Act continues by providing some examples: "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. The capacity of an AI system to infer transcends basic data processing by enabling learning, reasoning or modelling."

Complementarily, the Guidelines attempt to differentiate systems that have narrow inference capabilities, explaining that these systems have a "limited capacity to analyse patterns and adjust autonomously their output" (paragraph 41, GEC). Examples provided by the Commission include (paragraphs 42-51, GEC):

(i) Systems that improve mathematical optimization, or accelerate and approximate traditional optimization methods.
(ii) Systems that only exhibit basic data processing capabilities, which are: systems that follow predefined, and explicit instructions or operations, and which do not learn, reason or model at any stage of the system lifecycle; or systems for descriptive analysis, hypothesis testing and visualization.[2]
(iii) Systems based on classical heuristics, which are problem-solving techniques relying on experience-based methods to find approximate (not exact) solutions efficiently. Classical heuristics are not considered data-driven approaches and include rule-based approaches, pattern recognition, and trial-and-error strategies.
(iv) Systems that perform simple predictions, understood as systems whose performance can be achieved via a basic statistical learning rule. This includes, according to the Commission, systems that may be classified as relying on machine learning techniques.[3]

[2]This makes reference to paragraph 42, GEC, whose interpretation is not straightforward: "Systems used to improve mathematical optimisation or to accelerate and approximate traditional, well established optimisation methods, such as linear or logistic regression methods, fall outside the scope of the AI system definition." The Commission is likely not referring to linear and logistic regression as optimization methods, but rather to systems used to optimize or accelerate the computation of linear and logistic regression models. Therefore, it is likely not directly trying to exclude linear and logistic regression methods, but the methods used for optimization.

[3]This makes reference to paragraph 49, GEC, whose interpretation is not straightforward: "All machine-based systems whose performance can be achieved via a basic statistical learning rule, while technically may be classified as relying on machine learning approaches fall outside the scope of the AI system definition, due to its performance." Although the AI Act explicitly recognizes machine learning approaches within its scope (recital 12), the Guidance appears to introduce the possibility that some machine learning techniques may fall outside the scope of the AI Act. The Commission appears to base the exclusion of these systems on their performance. However, the performance criteria or threshold by which such systems would be excluded remains undefined.

3.3. Interaction and cumulative interpretation of constituents

Article 3(1) structures the constituents cumulatively (‘AI system’ means a machine-based system that … and … and ... and). However, as discussed above, several constituents are not decisive and have a limited exclusionary effect. The Commission clarifies that the seven constituents are not required to be continuously present throughout the AI system lifecycle (paragraph 10, GEC). The Commission acknowledges relationships between individual constituents, such as between autonomy and inference (paragraph 15, GEC), autonomy and adaptiveness (paragraph 22, GEC), inference and output generation (paragraph 26, GEC), and outputs and influence on the environment (paragraph 60, GEC). These relationships reflect the definition established in Article 3(1) of the AI Act. However, neither the AI Act nor the Commission, explains how these relationships, or the seven constituents collectively, should be interpreted when assessing the applicability of Article 3(1) of the AI Act.

4. Navigating the AI System Definition under the AI Act (NAIS) Framework, a framework for operationalizing the constituents of the AI system definition

In this section, the seven constituents of the AI system definition (European Commission 2025) were operationalized into a set of analytical questions intended to guide the assessment of the applicability of Article 3(1) of the AI Act (table 1, Appendix B). These questions constitute the “Navigating the AI System Definition under the AI Act” (NAIS) Framework. Each was formulated as a yes/no question followed by open questions that clarify the answer, this supports structured legal analysis rather than facilitating the deterministic classification of a system.

The formulation of the questions was primarily guided by the wording of the AI Act and, where necessary, further informed by the GEC. For constituents (2) autonomy and (5) inference, operationalization distinguishes between the development and deployment phases of the system lifecycle, in line with the Commission (paragraph 10, GEC). This distinction was introduced because autonomy and inference capabilities may change throughout the AI system lifecycle (paragraph 10, GEC). It was considered that constituent (4) exhibits objectives, may also change depending on the lifecycle stage. Because of their nature, constituents (1) machine-based, (3) adaptiveness, (6) outputs, and (7) influence, were considered only after deployment. Each operationalized constituent was subsequently applied to the SPIDeRR model, a logistic regression-based prediction system (table 1). The formulation of the questions was iterative and required consideration of the technical characteristics relevant to logistic regression systems in light of the AI system definition. Each question intends to elicit a yes/no answer, followed by a justification.

In summary (see table 1 for more details), the SPIDeRR logistic regression model constitutes a machine-based system (software only), with the explicit, human-predefined objective of estimating the presence or *de novo* diagnosis of a set of rheumatic diseases. For training, the variables and data were predefined by humans. During development, training was initiated by humans, during which the system exhibited a degree of operational autonomy characteristic of logistic regression training processes, and exhibited inference capabilities, generating the SPIDeRR model. During deployment, the system’s autonomy and inference capabilities may be considered very limited or absent, since the deployed model operates through a fixed mathematical structure (the model) applied to a small set of data entered manually by the user to generate the user predictions, without exhibiting adaptiveness. The underlying mathematical calculation is simple enough that in principle it could be performed manually by a human. For these reasons, it is possible to consider the system as having narrow inference capabilities, but it does not seem to match any of the examples provided by the Commission (paragraphs 42-51, GEC). The generated outputs may influence digital systems (such as healthcare software systems, e.g., electronic health records, if integrated into healthcare) and physical environments (such as human users, both patients and physicians, and healthcare pathways).

Table 1. NAIS Framework. Constituents, operational questions, and the SPIDeRR system as a use case.

Operational question(s)	Response [Y/N and explanation]
(1) The system is machine-based	
Was the system developed on machines* (hardware or software components), and does it run on machines?	Yes. It has only software components running on machines.
If yes, briefly describe the computational approach (e.g., mathematical model, or technique). If not, please explain. Does it exhibit computational capacities but is it implemented in another substrate?	The system is software only, implementing a logistic regression model.
*In this context, "machines" refer to computational systems including quantum computing.	

(2) The system is designed to operate with varying levels of autonomy Two phases are distinguished, development and deployment, because the level of autonomy of a system may differ across the system's lifecycle.	
Once deployed, does the system initiate or execute any operational processes* without human intervention?	No.
Please describe the processes carried out by the system once deployed and describe whether these processes are executed with or without human intervention.	The system carries out the following processes/functions during deployment: - Data acquisition: input data are manually entered by the user and are not autonomously acquired by the system. - Data processing, inferential processing, and output generation: user's data processing is initiated by the user, subsequently executed by the system using a previously trained and fixed logistic regression model. The system automatically generates a tailored output for the user.
Once deployed, is processing of (personal)ˆ data initiated only after an explicit action by a human user or a human operator? Please elaborate. If not, please describe if processing can be initiated autonomously by the system.	Yes. Data processing is carried out after data subjects provide explicit consent for processing.
During development, did the system initiate or execute any operational processes* without human intervention?	Yes.
Please describe the processes carried out by the system during development, and describe whether these processes are executed with or without human intervention	The system carried out the following processes/functions during development: - Data acquisition: input data are entered manually by users and are not autonomously acquired by the system. These data constituted the training dataset. - Data processing, inferential processing (including optimization), and output generation: once the training procedure was initiated by humans, the system computationally performed parameter estimation and coefficient optimization without direct human execution of the underlying mathematical operations. This is characteristic of a logistic regression algorithm.
During development, is processing of (personal)ˆ data initiated only after an explicit action by a human user or a human operator? Please elaborate. If not, please describe if processing can be initiated autonomously by the system.	Yes. Data subjects provided explicit consent to process their data for research purposes, and training was initiated by a researcher.
*Relevant processes may include data acquisition, data processing, inferential processing, output generation, action execution, monitoring, optimization, physical control, and interaction with users or other systems (appendix A provides examples of each category). ˆPersonal data as defined in the EU General Data Protection Regulation (GDPR): any information relating to an identified or identifiable natural person.	

(3) The system may exhibit adaptiveness after deployment	
Does the system modify its parameters, inferential processes, decision rules, or outputs after deployment based on new data or user interactions? Please justify.	No. The logistic regression model is fixed after training. The model applies predefined numerical values (coefficients) to variables relevant to the prediction, such as smoking status or pain. During deployment, the system adds together the values associated with the user's characteristics (e.g., if a user smokes) to generate a final probability score.

(4) The system exhibits explicit or implicit objectives	
Are the internal objectives (computational objectives) and intended use of the system explicitly specified (encoded) by developers, operators, or users? Specify the computational objective and the intended use. If there are differences between the development and deployment phases, please elaborate.	Yes. The model operates according to predefined objectives established by the developers (researchers). Computational objective: estimate the probability of a predefined outcome, specifically the presence or *de novo* diagnosis of a set of rheumatic diseases. Intended use: generate a probabilistic risk score and/or risk category for the user (patient and physician).
Does the system exhibit objectives that emerge implicitly from its design, training, optimization, or interaction with the environment? Please elaborate.	No. All objectives are explicit.

(5) The system infers Two phases are distinguished, development and deployment, because the inferential capabilities of a system may differ across the system's lifecycle.	
Once deployed, does the system derive outputs (e.g., predictions, recommendations, content, or decisions) from inputs or data through learning, reasoning, modeling, or other inferential processes*? Please explain. If the answer is yes, does the system perform inference beyond executing a previously established and fixed computational model? Is the computational model deterministic? Please elaborate.	No. The inference capabilities during deployment are very limited. During the deployment phase, the system does not optimize, derive new models, parameters, or inferential rules from incoming data. Instead, it executes a previously trained logistic regression model through a fixed mathematical procedure based on the variables entered by the user. The calculations performed by the model are such that they could be calculated by hand by a human. No.
During development, did the system derive models, parameters, or algorithms from inputs or data through learning, reasoning, modeling, or other inferential processes?* Please describe.	Yes. During development, the system derived the model coefficients from training data through statistical modeling and parameter optimization. The variables were not chosen by the system, or through a data-driven process, but were established by the researchers.
*The capacity to infer transcends basic data processing by enabling learning, reasoning or modeling (Recital 12, AI Act). Examples of systems that fall outside the scope of the AI system definition because of their limited capacity to analyze patterns and autonomously adapt their outputs can be found in paragraphs 40-51, GEC.	

(6) The system generates outputs	
Does the system produce predictions, recommendations, classifications, content, decisions or other outputs? Please specify.	Yes. The system generates a prediction, namely the presence or *de novo* diagnosis of a disease, presented as a risk category and/or score.

(7) The system influences physical or virtual environments	
Can the outputs of the system influence virtual (digital) or physical environments?* Please specify.	Yes. The system may influence virtual environments through its integration into healthcare software systems and may influence individuals and healthcare pathways through the risk-related outputs it generates.
*Article 3(1) of the AI Act does not provide examples of physical or virtual environments. However, paragraph 60, GEC refers to tangible physical objects (e.g., robotic arms), as well as virtual spaces such as digital spaces, data flows, and software ecosystems.	

5. Discussion

The NAIS Framework provides a systematic structure for analyzing the constituents of the AI system definition under Article 3(1) of the AI Act. The application of the NAIS Framework to the SPIDeRR system illustrates the interpretive difficulties associated with each constituent, some of which have been identified in theoretical analyzes [15, 3, 16]. Beyond the individual constituents, the application of the framework also exposes challenges in interpreting the AI system definition as a cumulative definition, as will be discussed in this section.

The SPIDeRR system clearly fulfills constituents (1) machine-based system, (4) objectives, (6) outputs, and (7) influence on physical or virtual environments. However, these constituents are generic and may equally be fulfilled by simple software systems traditionally not considered AI systems (e.g., rule-based spell checkers, tax calculation software, and appointment scheduling systems based entirely on human-defined rules). As a result, these constituents alone appear insufficient to meaningfully delimit the scope of Article 3(1), and therefore the applicability of the AI Act. This observation is particularly relevant in light of Recital 12 of the AI Act, which reflects the legislative intention that the definition distinguish AI systems from simpler traditional software systems or programming approaches. The SPIDeRR system clearly fails to satisfy constituent (3) because it does not exhibit adaptiveness.

The key question then becomes whether the SPIDeRR system fulfills constituents (2) autonomy and (5) inference. But it is in relation to these constituents that the greatest interpretive uncertainty emerges.

The AI Act and the GEC do not clearly specify how constituents (2) autonomy and (5) inference are satisfied in the context of AI systems, and how to weigh the different levels of autonomy and inference capabilities across the lifecycle. If basic inferential or autonomous processes are sufficient for qualification under Article 3(1), the scope of the AI system definition may extend to a wide range of simple statistical or computational systems. Again, this seems to create tension with Recital 12 of the AI Act, which aims to restrict the scope of the AI Act. Because the AI Act regulates systems placed on the market or put into service, one could argue that greater weight should be given to the post-deployment characteristics, rather than to characteristics that existed only during development. However, the position of the Commission, which clarified that the seven constituents are not required to be continuously present throughout the AI system lifecycle (paragraph 10, GEC), appears difficult to reconcile with this interpretation.

As discussed before, the uncertainty surrounding the interpretation of constituents (2) and (5) makes it difficult to conclusively determine whether SPIDeRR falls within the definition of an AI system under Article 3(1) of the AI Act. However, even if it could be conclusively established whether these constituents are fulfilled, from a regulatory perspective, it is unclear how to deal with systems that satisfy only some, and not all, constituents. For example, would the SPIDeRR system qualify as an AI system if it fulfills the generic constituents —(1) machine-based system, (4) objectives, (6) outputs, and (7) influence on environments— and (5) inference, but fails to satisfy (2) autonomy and (3) adaptiveness?

More broadly, the use of the NAIS Framework to analyze the SPIDeRR system demonstrates that the assessment of whether a system falls within the AI system definition requires a case-by-case approach, as neither the AI Act nor the GEC establish exhaustive technical thresholds or categorical criteria for determining whether a constituent is satisfied. This analysis demonstrates the difficulties of interpreting a definition composed of seven constituents in the absence of clear qualification criteria for each individual constituent and for their cumulative interpretation. Although the formulation of Article 3(1) may have been intended to provide flexibility to accommodate rapid technological developments (Recital 12, AI Act), and to preserve technological neutrality, it came at the expense of legal certainty for anyone seeking to determine whether a system falls within the scope of the AI Act. This uncertainty may incentivize overly cautious interpretations of the AI system definition, as already observed in practice by our academic research group. This perception is confirmed, for example, by a Dutch NGO that recommended that organizations "err on the side of caution" when determining whether systems fall within the scope of the AI Act [17].

Overly cautious compliance practices may result in disproportionate and unnecessary compliance burdens for policymakers, regulators, manufacturers, developers, researchers, and deployers, by treating simple statistical systems like AI systems despite the uncertainty surrounding their qualification under Article 3(1), or simply because of defensive compliance. This reduces legal certainty regarding the actual scope of the Regulation. It may also contribute to inconsistent interpretation and application of the AI Act across EU Member States, which may affect the prospect of achieving a digital single market [18]. If Member States or national competent authorities interpret the definition of an AI system differently, deployers operating across borders may face different regulatory obligations depending on the jurisdiction. But more fundamentally, overly cautious compliance practices may distance the AI Act, in practice, from its objective of establishing a proportionate, risk-based regulatory framework for AI systems that supports responsible innovation (Article 1(1), AI Act). Given this situation, there is a strong and urgent case for developing clearer guidelines, grounded not only in theoretical interpretation, but also in the practical application of the AI system definition.

The NAIS Framework presented here should therefore be understood as a first step toward operationalizing the AI system definition using the constituents identified by the Commission. Further refinement and validation through the analysis of additional systems will be necessary. Future work may also explore whether additional system characteristics, such as the level of risk posed, transparency, and explainability, should play a role in determining which systems fall within the scope of the AI Act. Such characteristics are currently not considered under Article 3(1), despite the AI Act's objective of establishing a proportionate, risk-based regulatory framework.

While this article was being drafted, the EU adopted the Omnibus VII legislative package, introducing targeted amendments to the AI Act as part of its simplification agenda [19]. Among other changes, the amendments introduce a mechanism to coordinate the AI Act with sector-specific legislation, allowing AI Act requirements to be limited where equivalent AI-specific obligations exist under sectoral legislation. Nevertheless, the question of the applicability of the definition of AI system remains relevant, even for systems that fall under sector-specific regulation, such as the MDR. That is because some of the obligations established in the AI Act for AI systems may be incorporated into sector-specific regulatory frameworks.

6. Conclusion

This article analyzed the AI system definition under Article 3(1) of the AI Act, and developed the NAIS Framework, a structured analytical framework for operationalizing the AI system definition. The NAIS Framework makes legal reasoning explicit, potentially serving as a structured and transparent framework for the case-by-case assessment of whether a system falls within the AI system definition. This may be valuable to policymakers, regulators, researchers, developers, providers and deployers, and in general, anyone in need of assessing whether a system falls within the AI system definition of the AI Act.

Applying the NAIS Framework to the SPIDeRR system demonstrated the limited discriminatory capacity of several constituents, highlighted important uncertainties regarding the interpretation of autonomy and inference, and exposed the practical consequences of the absence of guidance regarding the cumulative interpretation of the constituents. For this reason, it calls for clearer interpretative guidance. While the NAIS Framework and the analysis presented in this article do not resolve either the limited discriminatory capacity or the interpretive ambiguities of the AI system definition, they are intended as an initial step toward addressing these challenges. Their purpose is to contribute to legal scholarship by informing future authoritative interpretation and possible legislative refinement.

Future research should further refine and validate the NAIS Framework through its application to other systems representing different computational approaches. Comparative analysis of national guidance interpreting the AI system definition should also be undertaken to assess the consistency of interpretation across Member States. Future research should also examine whether the current criteria used to determine the scope of the AI Act remain appropriate, or whether additional or alternative considerations should be incorporated.

Acknowledgments

I would like to deeply thank T. de Graaf and R. Knevel for the insightful discussions that influenced this work.

This work is part of the SPIDeRR (Stratification of Patients using advanced Integrative modelling of Data Routinely acquired for diagnosing Rheumatic complaints) project, which has received funding from the European Health and Digital Executive Agency (HaDEA) [101080711].

Declaration on Generative AI

During the preparation of this work, the author used OpenAI GPT-5.5 and Grammarly in order to perform grammar and spelling checks. The author reviewed and edited the content as needed and assumes full responsibility for the publication's content.

References

[1] AI HLEG, Ethics Guidelines for Trustworthy AI. High-Level Expert Group on AI., Technical Report, European Commission, 2019.

[2] H. Ruschemeier, J. Bareis, Searching for Harmonised Rules: Understanding the Paradigms, Provisions, and Pressing Issues in the Final EU AI Act, in: R. Gsenger, M.-T. Sekwenz (Eds.), Digital Decade, Nomos Verlagsgesellschaft mbH & Co. KG, 2025, pp. 41–94. doi:10.5771/9783748943990-41.

[3] A. Rossi, F. Gennari, I. Fagioli, A. Mazzarini, F. Moncelli, D. Amram, S. Crea, A. Parziale, The AI system Definition under the AI Act, a New Nomen Rosae?, in: HHAI-WS 2025: Workshops at the Fourth International Conference on Hybrid Human-Artificial Intelligence (HHAI), June 9–13, 2025, Pisa, Italy, 2025.

[4] R. Knevel, J. Knitza, A. Hensvold, A. Circiumaru, T. Bruce, S. Evans, T. Maarseveen, M. Maurits, L. Beaart-van De Voorde, D. Simon, A. Kleyer, M. Johannesson, G. Schett, T. Huizinga, S. Svanteson, A. Lindfors, L. Klareskog, A. Catrina, Rheumatic?—A Digital Diagnostic Decision Support Tool for Individuals Suspecting Rheumatic Diseases: A Multicenter Pilot Validation Study, Frontiers in Medicine 9 (2022) 774945. doi:10.3389/fmed.2022.774945.

[5] K. Lundberg, L. Qin, C. Aulin, W. E. Van Spil, M. P. Maurits, R. Knevel, Population-based user-perceived experience of Rheumatic?: A novel digital symptom-checker in rheumatology, RMD Open 9 (2023) e002974. doi:10.1136/rmdopen-2022-002974.

[6] S. J. Russell, P. Norvig, Artificial Intelligence: A Modern Approach, Prentice Hall Series in Artificial Intelligence, third edition, global edition ed., Pearson, Boston Columbus Indianapolis, 2016.

[7] E. Bisong, Logistic Regression, Apress, Berkeley, CA, 2019, pp. 243–250. doi:10.1007/978-1-4842-4470-8_20.

[8] D. M. Burns, Artificial intelligence isn't, Canadian Medical Association Journal 192 (2020) E290–E290. doi:10.1503/cmaj.74722.

[9] ACCIS, Multi-Stakeholder Consultation for Commission Guidelines on the Application of the Definition of an AI System and the Prohibited AI Practices Established in the AI Act, Technical Report, ACCIS Data for credit, 2024.

[10] IAIS, Resolution of Public Consultation Comments on Draft Application Paper on the Supervision of Artificial Intelligence 18-11-24 to 17-02-25, Technical Report, IAIS, 2025.

[11] European Commission, Commission Guidelines on the Definition of an Artificial Intelligence System Established by Regulation (EU) 2024/1689 (AI Act). C(2025) 5053 Final, Brussels, 29 July 2025., Technical Report, 2025.

[12] S. Jakobi, K. Boy, M. Wagner, S. May, A. Temiz, A.-M. Liphardt, E. Araujo, L. Carmona, R. Knevel, G. Schett, J. Knitza, F. Muehlensiepen, H. Morf, Rheumatic? A diagnostic decision support tool for

individuals suspecting rheumatic diseases: Mixed-methods usability and acceptability study, BMC Rheumatology 9 (2025) 59. doi:10.1186/s41927-025-00507-w.
[13] OECD, OECD Framework for the Classification of AI Systems, Technical Report, 2022.
[14] Dutch Ministy Economic Affairs, AI Act Guide. Version 1.1, Technical Report, 2025.
[15] M. Almada, What is an Artificial Intelligence System, really? notes on the guidelines on Article 3(1) of the AI Act, Journal of AI Law and Regulation 2 (2025) 76–81. doi:10.21552/aire/2025/1/9.
[16] European Law Institute, Commission Guidelines on the Application of the Definition of an AI System and the Prohibited AI Practices Established in the AI Act. Respose of the European Law Institute., Technical Report, 2024.
[17] Algorithm Audit, Policy Briefing. Guidelines of the European Commission Blur the Interpretation of the AI System Definition, Technical Report, 2025.
[18] European Commission, Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions on Enabling the Digital Transformation of Health and Care in the Digital Single Market; empowering citizens and building a healthier society. COM/2018/233 final, 2018.
[19] Council of the EU, Artificial Intelligence: Council and Parliament agree to simplify and streamline rules. Press release. 7 May 2026 (2026).

Cited law

Dutch Civil Code – Book 7: Particular agreements (Articles 446–468: WGBO), https://wetten.overheid.nl/BWBR0005290/2020-01-01

Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices (MDR), http://data.europa.eu/eli/reg/2017/745/oj

Regulation (EU) 2023/988 of the European Parliament and of the Council of 10 May 2023 on general product safety (General Product Safety Regulation), http://data.europa.eu/eli/reg/2023/988/oj

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (AI Act), http://data.europa.eu/eli/reg/2024/1689/oj

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation, GDPR), http://data.europa.eu/eli/reg/2016/679/oj

Appendix

A. Systems processes (with examples)

The following list is illustrative rather than exhaustive. Not all systems perform all these processes or functions. The list provides examples of processes that systems may perform throughout their lifecycle and is intended to support the identification of processes relevant to the assessment of a system using the NAIS Framework.

Data acquisition: collecting sensor data, user inputs, images, text, or data from databases.

Data processing: cleaning, filtering, transformation, normalization, standardization, aggregation, feature extraction.

Optimization: parameter estimation, loss minimization, objective function optimization. Inferential processing: estimation, prediction, classification, ranking, or generation.

Output generation: recommendations, scores, labels, generated text, images, audio, or video.

Action execution: automatically triggering alerts, initiating or modifying workflows, scheduling appointments, contacting users or patients, or updating electronic health records.

Monitoring: monitoring data, system performance, or system behavior to detect conditions, changes, or anomalies.

Physical control: controlling robots, vehicles, industrial equipment, or other physical systems.

Communication: exchanging information with users, software systems, or AI systems.

B. Navigating the AI System Definition under the AI Act (NAIS) Framework

(1) The system is machine-based

Was the system developed on machines* (hardware or software components), and does it run on machines? If yes, briefly describe the computational approach (e.g., mathematical model, or technique). If not, please explain. Does it exhibit computational capacities but is it implemented in another substrate, such as biological or organic substrates?

**In this context, "machines" refer to computational systems including quantum computing.*

(2) The system is designed to operate with varying levels of autonomy

Two phases are distinguished, development and deployment, because the level of autonomy of a system may differ across the system's lifecycle.

Once **deployed**, does the system initiate or execute any operational processes* without human intervention? Please describe the processes carried out by the system once deployed and describe whether these processes are executed with or without human intervention. Once deployed, is processing of (personal)^ data initiated only after an explicit action by a human user or a human operator? Please elaborate. If not, please describe if processing can be initiated autonomously by the system.

During **development**, did the system initiate or execute any operational processes* without human intervention? Please describe the processes carried out by the system during development, and describe whether these processes are executed with or without human intervention.

During development, is processing of (personal)^ data initiated only after an explicit action by a human user or a human operator? Please elaborate. If not, please describe if processing can be initiated autonomously by the system.

**Relevant processes may include data acquisition, data processing, inferential processing, output generation, action execution, monitoring, optimization, physical control, and interaction with users or other systems (appendix 1 provides examples of each category).*

^Personal data as defined in the EU General Data Protection Regulation (GDPR): any information relating to an identified or identifiable natural person.

(3) The system may exhibit adaptiveness after deployment

Does the system modify its parameters, inferential processes, decision rules, or outputs after deployment based on new data or user interactions? Please justify.

(4) The system exhibits explicit or implicit objectives

Are the internal objectives (computational objectives) and intended use of the system explicitly specified (encoded) by developers, operators, or users? Specify the computational objective and the intended use. If there are differences between the development and deployment phases, please elaborate.

Does the system exhibit objectives that emerge implicitly from its design, training, optimization, or interaction with the environment? Please justify.

(5) The system infers

Two phases are distinguished, development and deployment, because the inferential capabilities of a system may differ across the system's lifecycle.

Once **deployed**, does the system derive outputs (e.g., predictions, recommendations, content, or decisions) from inputs or data through learning, reasoning, modeling, or other inferential processes*? Please explain. If the answer is yes, does the system perform inference beyond executing a previously established and fixed computational model? Is it the computational model deterministic? Please elaborate.

During **development**, did the system derive models, parameters, or algorithms from inputs or data through learning, reasoning, modeling, or other inferential processes?* Please describe.

**The capacity to infer transcends basic data processing by enabling learning, reasoning or modeling (Recital 12, AI Act). Examples of systems that fall outside the scope of the AI system definition because of their limited capacity to analyze patterns and autonomously adapt their outputs can be found in paragraphs 40-51, GEC.*

(6) The system generates outputs

Does the system produce predictions, recommendations, classifications, content, decisions or other outputs? Please specify.

(7) The system influences physical or virtual environments

Can the outputs of the system influence virtual (digital) or physical environments?* Please specify.

**Article 3(1) of the AI Act does not provide examples of physical or virtual environments. However, paragraph 60, GEC refers to tangible physical objects (e.g., robotic arms), as well as virtual spaces such as digital spaces, data flows, and software ecosystems.*