

EUROPEAN
COMMISSION

Brussels, 29.7.2025
C(2025) 5053 final

COMMUNICATION FROM THE COMMISSION

Commission Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689 (AI Act)

I. Purpose of the Guidelines

(1) Regulation (EU) 2024/1689 of the European Parliament and of the Council ('the AI Act')[1] entered into force on 1 August 2024. The AI Act lays down harmonised rules for the development, placing on the market, putting into service, and use of artificial intelligence ('AI') in the Union.[2] Its aim is to promote innovation in and the uptake of AI, while ensuring a high level of protection of health, safety, and fundamental rights in the Union, including democracy and the rule of law.

(2) The AI Act does not apply to all systems, but only to those systems that fulfil the definition of an 'AI system' within the meaning of Article 3(1) AI Act. The definition of an AI system is therefore key to understanding the scope of application of the AI Act.

(3) Article 96(1)(f) AI Act requires the Commission to develop guidelines on the application of the definition of an AI system as set out in Article 3(1) of that Act. By issuing these Guidelines, the Commission aims to assist providers and other relevant persons, including market and institutional stakeholders, in determining whether a system constitutes an AI system within the meaning of the AI Act, thereby facilitating the effective application and enforcement of that Act.

(4) The definition of an AI system entered into application on 2 February 2025[3], together with other provisions set out in Chapters I and II AI Act, notably Article 5 AI Act on prohibited AI practices. As the definition of an AI system is decisive to understanding the scope of the AI Act including the prohibited practices, the present Guidelines are adopted in parallel to Commission guidelines on prohibited artificial intelligence practices.

(5) These Guidelines take into account the outcome of a stakeholder consultation and the consultation of the European Artificial Intelligence Board.

(6) Considering the wide variety of AI systems, it is not possible to provide an exhaustive list of all potential AI systems in these Guidelines. This is in line with recital 12 AI Act, which clarifies that the notion of an 'AI system' should be clearly defined while providing 'the flexibility to accommodate the rapid technological developments in this field'. The definition of an AI system should not be applied mechanically; each system must be assessed based on its specific characteristics.

(7) The Guidelines are not binding. Any authoritative interpretation of the AI Act may ultimately only be given by the Court of Justice of the European Union (CJEU).

II. Objective and main elements of the AI system definition

(8) Article 3 (1) of the AI Act defines an AI system as follows:

[1] Regulation (EU) 2024/1689.
[2] Article 1 AI Act.
[3] Article 113, third paragraph, point (a).

"'AI system' means 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;"

(9) That definition comprises seven main elements: (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.

(10) The definition of an AI system adopts a lifecycle-based perspective encompassing two main phases: the pre-deployment or 'building' phase of the system and the post-deployment or 'use' phase of the system[4]. The seven elements set out in that definition are not required to be present continuously throughout both phases of that lifecycle. Instead, the definition acknowledges that specific elements may appear at one phase, but may not persist across both phases. This approach to define an AI system reflects the complexity and diversity of AI systems, ensuring that the definition aligns with the AI Act's objectives by accommodating a wide range of AI systems.

1. Machine-based system

(11) The term 'machine-based' refers to the fact that AI systems are developed with and run on machines. The term 'machine' can be understood to include both the hardware and software components that enable the AI system to function. The hardware components refer to the physical elements of the machine, such as processing units, memory, storage devices, networking units, and input/output interfaces, which provide the infrastructure for computation. The software components encompass computer code, instructions, programs, operating systems, and applications that handle how the hardware processes data and performs tasks.

(12) All AI systems are machine-based, since they require machines to enable their functioning, such as model training, data processing, predictive modelling and large-scale automated decision making. The entire lifecycle of advanced AI systems relies on machines that can include many hardware or software components. The element of 'machine-based' in the definition of AI system underlines the fact that AI systems must be computationally driven and based on machine operations.

(13) The term 'machine-based' covers a wide variety of computational systems. For example, the currently most advanced emerging quantum computing systems, which represent a significant departure from traditional computing systems, constitute machine-based systems, despite their unique operational principes and use of quantum-mechanical

[4] For overview of the AI system phases see the OECD (2024), "Explanatory memorandum on the updated OECD definition of an AI system", OECD Artificial Intelligence Papers, No. 8, OECD Publishing, Paris, https://doi.org/10.1787/623da898-en, p.7.

phenomena, as do biological or organic systems so long as they provide computational capacity.

2. Autonomy

(14) The second element of the definition refers to the system being 'designed to operate with varying levels of autonomy'. Recital 12 of the AI Act clarifies that the terms 'varying levels of autonomy' mean that AI systems are designed to operate with 'some degree of independence of actions from human involvement and of capabilities to operate without human intervention'.

(15) The notions of autonomy and inference go hand in hand: the inference capacity of an AI system (i.e., its capacity to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments) is key to bring about its autonomy.

(16) Central to the concept of autonomy is 'human involvement' and 'human intervention' and thus human-machine interaction. At one extreme of possible human-machine interaction are systems which are designed to perform all tasks though manually operated functions. At the other extreme are systems that are capable to operate without any human involvement or intervention, i.e. fully autonomously.

(17) The reference to 'some degree of independence of action' in recital 12 AI Act excludes systems that are designed to operate solely with full manual human involvement and intervention. Human involvement and human intervention can be either direct, e.g. through manual controls, or indirect, e.g. though automated systems-based controls which allow humans to delegate or supervise system operations.

(18) For example, a system that requires manually provided inputs to generate an output by itself is a system with 'some degree of independence of action', because the system is designed with the capability to generate an output without this output being manually controlled, or explicitly and exactly specified by a human. Likewise, an expert system following a delegation of process automation by humans that is capable, based on input provided by a human, to produce an output on its own such as a recommendation is a system with 'some degree of independence of action'.

(19) The reference in the definition of an AI system in Article 3(1) AI Act to 'machine-based system that is designed to operate with the varying levels of autonomy' underlines the ability of the system to interact with its external environment, rather than a choice of a specific technique, such as machine learning, or model architecture for the development of the system.

(20) Therefore, the level of autonomy is a necessary condition to determine whether a system qualifies as an AI system. All systems that are designed to operate with some reasonable degree of independence of actions fulfil the condition of autonomy in the definition of an AI system.

(21) Systems that have the capability to operate with limited or no human intervention in specific use contexts, such as in the high-risk areas identified in Annex I and Annex III AI Act, may, under certain conditions, trigger additional potential risks and human oversight considerations. The level of autonomy is an important consideration for a provider when devising, for example, the system's human oversight or risk mitigation measures in the context of the intended purpose of a system.

3. Adaptiveness

(22) The third element of the definition in Article 3(1) AI Act is that the system 'may exhibit adaptiveness after deployment'. The concepts of autonomy and adaptiveness are two distinct but closely related concepts. They are often discussed together but they represent different dimensions of an AI system's functionality. Recital 12 AI Act clarifies that 'adaptiveness' refers to self-learning capabilities, allowing the behaviour of the system to change while in use. The new behaviour of the adapted system may produce different results from the previous system for the same inputs.

(23) The use of the term 'may' in relation to this element of the definition indicates that a system may, but does not necessarily have to, possess adaptiveness or self-learning capabilities after deployment to constitute an AI system. Accordingly, a system's ability to automatically learn, discover new patterns, or identify relationships in the data beyond what it was initially trained on is a facultative and thus not a decisive condition for determining whether the system qualifies as an AI system.

4. AI system objectives

(24) The fourth element of the definition is AI system objectives. AI systems are designed to operate according to one or more objectives. The objectives of the system may be explicitly or implicitly defined. **Explicit objectives** refer to clearly stated goals that are directly encoded by the developer into the system. For example, they may be specified as the optimisation of some cost function, a probability, or a cumulative reward. **Implicit objectives** refer to goals that are not explicitly stated but may be deduced from the behaviour or underlying assumptions of the system. These objectives may arise from the training data or from the interaction of the AI system with its environment.

(25) Recital 12 AI Act clarifies that, 'the objectives of the AI system may be different from the intended purpose of the AI system in a specific context'. The objectives of an AI system are internal to the system, referring to the goals of the tasks to be performed and their results. For instance, a corporate virtual AI assistant system may have objectives to answer user questions on a set of documents with high accuracy in and low rate of failures. In contrast, the **intended purpose** is externally oriented and includes the context in which the system is designed to be deployed and how it must be operated. Indeed, according to Article 3(12) AI Act, the intended purpose of an AI system refers to the 'use

for which an AI system is intended by the provider'. For example, in the case of a corporate virtual AI assistant system, the intended purpose might be to assist a certain department of a company to carry out certain tasks. This might require that the documents that the virtual assistant uses comply with certain requirements (e.g. length, formatting) and that the user questions are limited to the domain in which the system is intended to operate. This intended purpose is fulfilled not only through the system's internal operation to achieve its objectives, but also through other factors, such as the integration of the system into a broader customer service workflow, the data that is used by the system, or instructions for use.

5. Inferencing how to generate outputs using AI techniques

(26) The fifth element of an AI system is that it must be able to infer, from the input it receives, how to generate outputs. Recital 12 AI Act clarifies that "[a] key characteristic of AI systems is their capability to infer." As further explained in that recital, AI systems should be distinguished 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." This capability to infer is therefore a key, indispensable condition that distinguishes AI systems from other types of systems.

(27) Recital 12 also explains that '[t]his capability to infer refers to the process of obtaining the outputs, such as predictions, content, recommendations, or decisions, which can influence physical and virtual environments, and to a capability of AI systems to derive models or algorithms, or both, from inputs or data.' This understanding of the concept of 'inference' does not contradict the ISO/IEC 22989 standard, which defines inference 'as reasoning by which conclusions are derived from known premises' and this standard includes an AI-specific note stating: '[i]n AI, a premise is either a fact, a rule, a model, a feature or raw data." [5].

(28) The 'process of obtaining the outputs, such as predictions, content, recommendations, or decisions, which can influence physical and virtual environments', refers to the ability of the AI system, predominantly in the 'use phase', to generate outputs based on inputs. A 'capability of AI systems to derive models or algorithms, or both, from inputs or data' refers primarily, but is not limited to, the 'building phase' of the system and underlines the relevance of the techniques used for building a system.

(29) The terms 'infer how to', used in Article 3(1) and clarified in recital 12 AI Act, is broader than, and not limited only to, a narrow understanding of the concept of inference as an ability of a system to derive outputs from given inputs, and thus infer the result. Accordingly, the formulation used in Article 3(1) AI Act, i.e. 'infers, how to generate outputs', should be understood as referring to the building phase, whereby a system derives outputs through AI techniques enabling inferencing.

[5] ISO/IEC 22989:2022, Information technology — Artificial intelligence — Artificial intelligence concepts and terminology.

5.1. AI techniques that enable inference

(30) Focusing specifically on the building phase of the AI system, recital 12 AI Act further clarifies that '[t]he 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.' Those techniques should be understood as 'AI techniques'.

(31) This clarification explicitly underlines that the concept of 'inference' should be understood in a broader sense as encompassing the 'building phase' of the AI system. Recital 12 AI Act then provides further guidance on techniques that enable this ability of an AI system to infer how to generate outputs. Accordingly, the techniques that may be used to enable inference 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.'

(32) The first category of AI techniques mentioned in recital 12 AI Act is '**machine learning approaches** that learn from data how to achieve certain objectives'. That category includes a large variety of approaches enabling a system to 'learn', such as supervised learning, unsupervised learning, self-supervised learning and reinforcement learning.

(33) In the case of **supervised learning**, the AI system learns from annotations (labelled data), whereby the input data is paired with the correct output. The system uses those annotations to learn a mapping from inputs to outputs and then generalises this to new, unseen data. An AI-enabled e-mail spam detection system is an example of a supervised learning system. During its building phase, the system is trained on a dataset containing emails that humans have labelled as 'spam' or 'not spam' to learn patterns from the features of the labelled e-mails. Once trained and in use, the system can analyse new e-mails and classify them as spam or not spam based on the patterns it has learned from the labelled data.

(34) Other examples of AI systems based on supervised learning include image classification systems trained on a dataset of images, whereby each image is labelled with a set of labels (e.g. objects such as cars), medical device diagnostic systems trained on medical imaging labelled by human experts, and fraud detection systems that are trained on labelled transaction data.

(35) In the case of **unsupervised learning**, the AI system learns from data that has not been labelled. The model is trained on data without any predefined labels or outputs. Using different techniques, such as clustering, dimensionality reduction, association rule learning, anomality detection, or generative models, the system is trained to find patters, structures or relationships in the data without explicit guidance on what the outcome should be. AI systems used for drug discovery by pharmaceutical companies is an

example of unsupervised learning. AI systems use unsupervised learning (e.g. clustering or anomality detection) to group chemical compounds and predict potential new treatments for diseases based on their similarities to existing drugs.

(36) **Self-supervised learning** is a subcategory of unsupervised learning, whereby the AI system learns from unlabelled data in a supervised fashion, using the data itself to create its own labels or objectives. AI systems based on self-supervised learning use various techniques, such as auto-encoders, generative adversarial networks, or contrastive learning. An image recognition system that learns to recognise objects by predicting missing pixels in an image is an example of an AI system based on self-supervised learning. Other examples include language models that learn to predict the next token in a sentence or speech recognition systems that learn to recognise spoken words by predicting the next acoustic feature in an audio signal.

(37) AI systems based on **reinforcement learning** learn from data collected from their own experience through a 'reward' function. Unlike AI systems that learn from labelled data (supervised learning) or that learn from patterns (unsupervised learning), AI systems based on reinforcement learning learn from experience. The system is not given explicit labels but instead learns by trial and error, refining its strategy based on the feedback it gets from the environment. An AI-enabled robot arm that can perform tasks like grasping objects is an example of an AI system based on reinforcement learning. Reinforcement learning can be also used, for example, to optimise personalised content recommendations in search engines and the performance of autonomous vehicles.

(38) **Deep learning** is a subset of machine learning that utilises layered architectures (neural networks) for representation learning. AI systems based on deep learning can automatically learn features from raw data, eliminating the need for manual feature engineering. Due to the number of layers and parameters, AI systems based on deep learning typically require large amounts of data to train, but can learn to recognise patterns and make predictions with high accuracy when given sufficient data. AI systems based on deep learning are widely used, and it is a technology behind many recent breakthroughs in AI.

(39) In addition to various machine learning approaches discussed above, the second category of techniques mentioned in recital 12 AI Act are '**logic- and knowledge-based approaches** that infer from encoded knowledge or symbolic representation of the task to be solved'. Instead of learning from data, these AI systems learn from knowledge including rules, facts and relationships encoded by human experts. Based on the human experts encoded knowledge, these systems can 'reason' via deductive or inductive engines or using operations such as sorting, searching, matching, chaining. By using logical inference to draw conclusions, such systems apply formal logic, predefined rules or ontologies to new situations. Logic- and knowledge-based approaches include for instance, knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning, expert systems and search and optimisation methods. For example, classical language processing models based on grammatical knowledge and logical semantics rely on the structure of language,

identifying the syntactical and grammatical components of sentences to extract the meaning of a given text. Another prominent example of AI systems based on logic and knowledge-based approaches are early generation expert systems intended for medical diagnosis, which are developed by encoding knowledge of a range of medical experts and which are intended to draw conclusions from a set of symptoms of a given patient.

5.2. Systems outside the scope of the AI system definition

(40) Recital 12 also explains that the AI system definition should distinguish AI systems 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."

(41) Some systems have the capacity to infer in a narrow manner but may nevertheless fall outside of the scope of the AI system definition because of their limited capacity to analyse patterns and adjust autonomously their output. Such systems may include:

Systems for improving mathematical optimization

(42) 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. This is because, while those models have the capacity to infer, they do not transcend 'basic data processing'. An indication that a system does not transcend basic data processing could be that it has been used in consolidated manner for many years[6]. This includes, for example, machine learning-based models that approximate functions or parameters in optimization problems while maintaining performance. The systems aim to improve the efficiency of optimisation algorithms used in computational problems. For example, they help to speed up optimisation tasks by providing learned approximations, heuristics, or search strategies.

(43) For example, physics-based systems may use machine learning techniques to improve computational performance, accelerating traditional physics-based simulations or estimating parameters, that are then fed into the established physics models. These systems would fall outside the scope of the AI system definition. In this example, machine learning models approximate complex atmospheric processes, such as cloud microphysics or turbulence, enabling faster and more computationally efficient forecasts.

(44) Another example of a system that falls outside the scope of the definition is a satellite telecommunication system to optimize bandwidth allocation and resource management. In satellite communication, traditional optimization methods may struggle with real-time demands of network traffic, especially when adjusting for varying levels of user demand across different regions. Machine learning models, for instance, can be used to predict

[6] In any case, the systems that are already placed on the market or put into service before 2 August 2026 benefit from 'grandfathering' clause foreseen in Article 111(2) AI Act.

network traffic and optimize the allocation of resources like power and bandwidth to satellite transponders, having similar performance to established methods in the field.

(45) Whilst these systems may incorporate automatic self-adjustments, these adjustments are addressed at optimising the functioning of the systems by improving its computational performance rather than, for example, at permitting adjustments of their decision making models in an intelligent way. Under these conditions they may be excluded from the AI system definition.

Basic data processing

(46) Basic data processing system refers to a system that follows predefined, explicit instructions or operations. These systems are developed and deployed to execute tasks based on manual inputs or rules, without any 'learning, reasoning or modelling' at any stage of the system lifecycle. They operate based on fixed human-programmed rules, without using AI techniques, such as machine learning or logic-based inference, to generate outputs. These basic data processing systems include, for example, database management systems used to sort or filter data based on specific criteria (e.g. 'find all customers who purchased a specific product in the last month'), standard spreadsheet software applications which do not incorporate AI enabled functionalities, and software that calculates a population average from a survey that is later exploited in a general context.

(47) Also systems that solely intended for descriptive analysis, hypothesis testing, and visualisation, fall outside the definition of an AI system. For instance, in software for sales report visualisation, statistical methods can be used to create a sales dashboard that shows total sales, average sales per region and sales trends over time. With the help of statistical methods, those data can be summarised and visualised in charts and graphs. However, the system does not recommend how to improve sales or which products to promote. Another example is a software system that applies statistical techniques to opinion polls or survey data to determine their validity, reliability, correlation, and statistical significance. Such systems do not 'learn, reason or model', they simply present data in an informative way.

Systems based on classical heuristics

(48) Classical heuristics are problem-solving techniques that rely on experience-based methods to find approximate solutions efficiently. Heuristics techniques are commonly used in programming situations where finding an exact solution is impractical due to time or resource constraints. Classical heuristics typically involve rule-based approaches, pattern recognition, or trial-and-error strategies rather than data-driven learning. Unlike modern machine learning systems, which adjust their models based on input-output relationships, classical heuristic systems apply predefined rules or algorithms to derive solutions. For instance, a chess program using a minimax algorithm with heuristic evaluation functions can assess board positions without requiring prior learning from data. While effective in many applications, heuristic methods may lack adaptability and generalization compared to AI systems that learn from experience.

Simple prediction systems

(49) 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.

(50) For instance, in financial forecasting (basic benchmarking) such machine-based systems may be used to predict future stock prices by using an estimator with the 'mean' strategy to establish a baseline prediction (e.g., always predicting the historical average price). Such basic benchmarking methods help to assess whether more advanced machine learning models could add value. Another example is using the average temperature of last week for predicting tomorrow's temperature. This baseline system solely estimates averages, but it is not achieving the performance of more complex time-series forecasting systems that would require more sophisticated models.

(51) Static estimation systems, such as customer support response time system that are based on static estimation to predict the mean resolution time from the past data and trivial predictors such as demand forecasting for a store to predict how many items of a product the store will sell each day are other examples, that help to establish a baseline or a benchmark, e.g. by predicting average or mean.

6. Outputs that can influence physical or virtual environments

(52) The sixth element of the AI system definition in Article 3(1) AI Act is that the system infers 'how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments'. The ability of a system, to generate outputs, such as predictions, content, and recommendations, based on inputs it receives and using machine learning and logic and knowledge-based approaches, is fundamental to what AI systems do and what distinguishes those systems from other forms of software. The capacity to generate outputs and the type of output the system can generate is central to understanding the functionality and impact of an AI system.

(53) Outputs of AI systems belong to four broad categories listed in Article 3(1) AI Act: predictions, content, recommendations, and decisions. Each category differs in its level of human involvement.

(54) **Predictions** are one of the most common outputs that AI system produce and that require the least human involvement. A prediction is an estimate about an unknown value (the output) from known values supplied to the system (the input). Software systems have been used for decades to generate predictions. AI systems using machine learning are capable of generating predictions that uncover complex patterns in data and make accurate predictions in highly dynamic and complex environments.

(55) For example, AI systems deployed in self-driving cars are designed to make real-time predictions in an extremely complex and dynamic environment, with multiple types of agents and interactions, and a practically infinite number of possible situations, and to

take decisions to adjust their behaviour accordingly. Non-AI systems, typically based on historical data, scientific data or predefined rules, such as certain non-AI medical device expert systems, are not capable of dealing with such a degree of complexity. Similarly, AI systems for energy consumption are designed to estimate energy consumption by analysing data from smart meters, weather forecasts and behavioural patterns on consumers. By relying on machine learning approaches, an AI system is designed to find complex correlations between these variables to make more accurate energy consumption predictions.

(56) **Content** refers to the generation of new material by an AI system. This may include text, images, videos, music and other forms of output. There is an increasing number of AI systems that use machine learning models (for example based on Generative Pre-trained Transformer (GPT) technologies) to generate content. Although content, as a category of output, may be understood from a technical perspective in terms of a sequence of 'predictions' or 'decisions', due to the prevalence of this output in generative AI systems, it is listed in recital 12 AI Act as a separate category of output.

(57) **Recommendations** refer to suggestions for specific actions, products, or services to users based on their preferences, behaviours, or other data inputs. Similarly to predictions, both AI-based and non-AI-based systems can be designed to generate recommendations. AI-based recommendation systems, for example, can leverage large-scale data, adapt to user behaviour in real-time, provide highly personalised recommendations, and scale efficiently as the dataset grows, the functionalities that non-AI systems that rely on static, rule-based mechanisms and limited data, rarely possess. In other cases, recommendations refer to potential decisions, such as a candidate to hire in a recruitment system, which will be evaluated by humans. If these recommendations are automatically applied, they become decisions.

(58) **Decisions** refer to conclusions or choices made by a system. An AI system that outputs a decision automates processes that are traditionally handled by human judgement. Such a system implies a fully automated process whereby a certain outcome is produced in the environment surrounding the system without any human intervention.

(59) In summary, AI systems, including systems based on machine learning approaches and logic or knowledge-based systems, differ from non-AI systems in their ability to generate outputs like predictions, content, recommendation, and decisions in that they can handle complex relationships and patterns in data. AI systems can generally generate more nuanced outputs than other systems, for example, by leveraging patterns learned during training or by using expert-defined rules to make decisions, offering more sophisticated reasoning in structured environments.

7. Interaction with the environment

(60) The seventh element of the definition of an AI system is that system's outputs **'can influence physical or virtual environments'**. That element should be understood to emphasise the fact that AI systems are not passive, but actively impact the environments

in which they are deployed. Reference to 'physical or virtual environments' indicates that the influence of an AI system may be both to tangible, physical objects (e.g. robot arm) and to virtual environments, including digital spaces, data flows, and software ecosystems.

III. Concluding remarks

(61) The definition of an AI system encompasses a wide spectrum of systems. The determination of whether a software system is an AI system should be based on the specific architecture and functionality of a given system and should take into consideration the seven elements of the definition laid down in Article 3(1) AI Act.

(62) No automatic determination or exhaustive lists of systems that either fall within or outside the definition of an AI system are possible.

(63) Only certain AI systems are subject to regulatory obligations and oversight under the AI Act. The AI Act's risk-based approach means that only those systems giving rise to the most significant risks to fundamental rights and freedoms will be subject to its prohibitions laid down in Article 5 AI Act, its regulatory regime for high-risk AI systems covered by Article 6 AI Act and its transparency requirements for a limited number of pre-defined AI systems laid down in Article 50 AI Act. The vast majority of systems, even if they qualify as AI systems within the meaning of Article 3(1) AI Act, will not be subject to any regulatory requirements under the AI Act.

(64) The AI Act also applies to general-purpose AI models, which are regulated in Chapter V of the AI Act. The analysis on the differences between AI systems and general-purpose AI models is outside the scope of these Guidelines.

CERTIFIED COPY
For the Secretary-General

Martine DEPREZ
Director
Decision-making & Collegiality
EUROPEAN COMMISSION