

The AI system Definition under the AI Act, a New Nomen Rosae?

Arianna Rossi^{1,*}, Francesca Gennari¹, Ilaria Fagioli^{2,3}, Alessandro Mazzarini^{2,3},
Fabrizio Moncelli^{2,3}, Denise Amram^{1,4}, Simona Crea^{2,3} and Andrea Parziale^{1,4}

¹LIDERLab, DIRPOLIS Institute, Sant'Anna School of Advanced Studies, Piazza dei Martiri della Libertà 33, 56127 Pisa, Italy

²WRLab, The BioRobotics Institute, Sant'Anna School of Advanced Studies, Viale Rinaldo Piaggio 34, 56025 Pontedera, Italy

³Department of Excellence in Robotics & AI, Scuola Superiore Sant'Anna, Piazza dei Martiri della Libertà 33, 56127 Pisa, Italy

⁴L'EMbeDS Department of Excellence, Piazza dei Martiri della Libertà 33, 56127 Pisa, Italy

⁴Health Science Interdisciplinary Center, Sant'Anna School of Advanced Studies, Piazza dei Martiri della Libertà 33, 56127 Pisa, Italy

Abstract

This short paper aims to establish how to apply the 'AI system' definition provided in the AI Act's Article 3(1) and further clarified in the dedicated European Commission's Guidelines in practice. Thanks to an interdisciplinary collaboration between legal scholars and bioengineers, we identify alignments and discrepancies between the legal definitions and the actual use of the same terms in the engineering community. We hence define a commented comparative interdisciplinary thesaurus of key terms. We further discuss the implications that terminological and interpretative issues may have on legal certainty and the compliance activities of developers of AI systems.

Keywords

AI Act, biorobotics, robotic prostheses, AI safety, law and tech

1. Introduction

The EU Regulation 2024/1689 on Artificial Intelligence (hereinafter: AI Act) constitutes a landmark legislation worldwide for the governance of AI. Although it aims to bring legal certainty to the development and use of AI systems in all contexts of human life, its approval has raised interpretative issues regarding the qualification of automated systems as AI systems within the meaning of this European regulation. Even though the AI Act does not apply to 'AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development' (Article 2(6)), many AI projects envision the real-world use and/or the commercialization of their results: bioengineering and biomedical research projects, for instance, almost always intend to have a concrete impact outside of research laboratories. AI developers and deployers thus need to understand whether and how the provisions of the AI Act apply to the systems they create or use. For example, in the case of a high-risk system, specific requirements (e.g., in terms of system design and documentation) need to be considered and implemented early on, since ignoring them would jeopardize later deployment opportunities [1].

This short paper illustrates the preliminary results that have emerged from an interdisciplinary collaboration between a team of legal scholars and a team of bioengineering researchers of the BRIEF

HHAI-WS 2025: Workshops at the Fourth International Conference on Hybrid Human-Artificial Intelligence (HHAI), June 9–13, 2025, Pisa, Italy

*Corresponding author.

✉ arianna.rossi@santannapisa.it (A. Rossi); francesca.gennari@santannapisa.it (F. Gennari); ilaria.fagioli@santannapisa.it (I. Fagioli); alessandro.mazzarini@santannapisa.it (A. Mazzarini); fabrizio.moncelli@santannapisa.it (F. Moncelli); denise.amram@santannapisa.it (D. Amram); simona.crea@santannapisa.it (S. Crea); andrea.parziale@santannapisa.it (A. Parziale)

ORCID 0000-0002-4199-5898 (A. Rossi); 0000-0003-0847-8466 (F. Gennari); 0000-0002-6379-7876 (I. Fagioli); 0000-0002-4524-7029 (A. Mazzarini); 0000-0001-7521-3148 (F. Moncelli); 0000-0002-4627-7502 (D. Amram); 0000-0001-9833-4401 (S. Crea); 0000-0002-3091-0705 (A. Parziale)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

project.¹ who design robotic prostheses. This use case is of particular interest as it constitutes an advanced medical device designed to replace a missing limb and restore motor functions through the integration of mechanical components and intelligent control systems [2] (see Sec. 2). First, we critically examine the notion of "AI system" provided in the AI Act in Article 3(1) and further clarified in Recital 12. This is key, because understanding if a robotic prosthesis qualifies as an AI system as defined in the AI Act determines whether its developers and deployers should respect the applicable obligations and requirements, depending on its level of risk. The cited Recital 12 states that "[t]he notion of 'AI system' in this Regulation should be clearly defined and should be closely aligned with the work of international organizations 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." With the intent to offer valuable elements for the interpretation of such crucial definition, in February 2025 the AI Office of the EU Commission published the *Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689 (AI Act)* [3].

The guidelines were first analyzed by the legal experts of the team who formulated interpretative hypotheses. Such hypotheses were then confirmed or reviewed by the bioengineers during iterative workshops in March - April 2025 (ca. 10 hours) to identify potential discrepancies in the use of terminology by the AI Office compared to the practices of the engineering community. Together, we determined the notions that are critical to determine the inclusion or exclusion of the robotic prosthesis within the boundaries of the AI system definition and sought to clarify them, by engaging in critical discussions across disciplinary domains. This short article presents the preliminary findings of such a fruitful collaboration in the form of a commented comparative vocabulary of the terminology used in the AI Act and in the EC's guidelines with technical handbook definitions, as summarized in Table 1. As a result, we argue that, even though some explanations of the guidelines may ease the application of the regulation, the use of certain terms complicates it further, thereby exacerbating legal uncertainty.

2. The use case: robotic prosthesis

In the framework of the BRIEF project, we considered as a use case scenario the robotic prostheses developed by the Sant'Anna School of Advanced Studies, i.e., a robotic knee [4] and a robotic ankle [5]. Each prosthesis contains an electric motor, batteries, and sensors to monitor the state of the device and the motion of the user's residual limb. Powered prostheses typically rely on a three-layered control architecture designed to mimic the coordinated movement of a biological limb. These control architectures are structured into multiple levels, each with a specific role in transforming user intent into mechanical actions. This hierarchical framework allows the prosthesis to interpret the user's needs, generate appropriate motor commands, and execute movements in real time [6]. In this specific use case, the low-level layer is a proportional-integrative-derivative (PID) controller that directly adjusts the motor's current based on the error between the measured and desired output torque of the device. The middle-layer computes in real-time the torque command required to modulate the impedance of the robotic joints, using a data-driven model trained offline from healthy people's data. The high-level layer identifies the specific locomotion task and sends relevant information to the middle layer. At the current stage of development, task selection is performed manually by a human operator, but future developments aim to automate this process.

3. When is an AI system...an AI system?

3.1. On the definition of AI system

An 'AI system' is defined as 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, contents,

¹"Biorobotics Research and Innovation Engineering Facilities" (BRIEF). Available at: <https://biorob-hub.eu/>

recommendations, or decisions that can influence physical or virtual environments” (Article 3(1)). Not all of these elements are required to be present continuously throughout the product life cycle. They may be present during the pre-deployment (or building) phase but be absent in the post-deployment (or use) phase, or vice versa [3].

Following the guidelines, we do not doubt that the robotic prosthesis is a machine-based system (i.e., AI systems “are developed with and run on machines” p. 2) that has objectives (i.e., it has goals for the tasks it needs to perform) and that influences a physical environment (i.e., it “actively impact[s] the environments in which [it is] deployed” p. 11-12). In the following, we analyze the other elements of the definition.

3.2. “Varying levels of autonomy”

Autonomy refers to the ability of AI systems to “have some degree of independence of actions from human involvement and of capabilities to operate without human intervention” (Recital 12). The guidelines further explain that human involvement can be either direct (e.g., manual control) or indirect (e.g., automated system-based control), and that a system is autonomous when it is “designed to operate with some reasonable degree of independence” (p. 3), without specifying what would amount to such a reasonable level. Lastly, the guidelines state that “independence of action” refers to the system’s capacity to generate an output “on its own”, without manual control or an explicit and exact specification by a human being.

Robotic lower limb prostheses exhibit a certain level of autonomy, which is consistent with the definition. However, the notion of autonomy is context-dependent. From the users’ perspective, the system remains inherently dependent on their movement: the prosthesis generally does not initiate motion independently but reacts to the user’s residual limb activity and intent. From the engineer’s standpoint, manual intervention may be required at the high-level control, such as when selecting a specific locomotion mode. Once this task is set, however, the prosthesis can execute low- and middle-level control commands autonomously, without requiring further human input. Thus, each component of an AI system may be characterized by a different level of autonomy. Nevertheless, these factors must all be considered when determining whether an AI system is autonomous.

3.3. “May exhibit adaptiveness”

Following Recital 12 of the AI Act, the guidelines clarify that adaptiveness “refers to self-learning capabilities, allowing the system to change while in use” [3, p. 4]. Further, they mention that adaptiveness may only pertain to the pre-deployment phase and is facultative in the post-deployment phase. We question whether adaptiveness always implies self-learning capabilities, even though the guidelines in the following paragraph mention them disjunctively: “the term ‘may’ [...] indicates that a system may, but does not necessarily have to, possess adaptiveness *or* self-learning capabilities after deployment” (p.4) [emphasis added].

The concept of adaptiveness in powered prostheses is often used broadly and there is no universally accepted definition. In general terms, a system is considered adaptive if it can modify its behavior in response to changes in the user or environment [7]. However, adaptiveness does not necessarily imply self-learning. For example, after deployment, the high-level controller may recognize the current locomotion task (e.g., level walking, stair ascent) and send this information to the middle-level controller. In turn, the middle-level controller adjusts its parameters, such as switching to a different set of control parameters that correspond to the identified task. In this case, the system adapts its behavior but does not generate new knowledge, nor does it update its internal models based on new data. Thus, it is adaptive, but it is not self-learning. Real self-learning systems, on the contrary, continuously update or refine their models based on real-time feedback or user-specific performance, even after deployment [8]. Therefore, while all self-learning systems are adaptive, not all adaptive systems are self-learning.

3.4. “Infers, from the input it receives, how to generate outputs”

Recital 12 specifies that 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.” Leaving aside the fact that the terms “simpler” and “traditional” are not unequivocally codified in this context, defining what falls and what does not fall under this definition of inference is critical: the definition risks being over-comprehensive, since most, if not all, control algorithms are based on inferences.

3.4.1. Rule-based systems

First, a clear dividing line is the contrast between AI systems and rule-based systems. The guidelines state that their definition does not contradict the ISO/IEC 22989’s [9] definition of inference: “reasoning by which conclusions are derived from known premises” such as “a fact, *a rule*, a model, a feature or raw data”. We posit that the AI Office intended to specify that the AI Act’s definition does not exclude rules in the context of known premises, but systems that operate solely based on expert-defined rules (i.e., rule-based) are. Hence, an algorithm that detects when the prosthesis is in contact with the ground based on a signal above a user-set threshold would be considered rule-based and would not fall under the AI Act’s definition.

3.4.2. Inference: obtaining outputs + deriving models or algorithms from inputs or data

The guidelines further explain that the ability to infer refers to:

1. “the process of obtaining the outputs” (Recital 12), which corresponds to the ability “to generate outputs based on inputs” [3, p. 5], which, according to the AI Office, is mainly tied to the use phase; and
2. the “capability of AI systems to derive models or algorithms, or both, from inputs or data” (Recital 12), which “underlines the relevance of the techniques used for building a system” [3, p. 5] and refers primarily to the building phase.

These two points underscore the difference between the use phase and the building phase. In the first case, the intention may be to emphasize the otherness of AI systems that can generate outputs from unknown data or new situations during deployment, unlike rule-based systems. In the second case, the explicit reference to models or algorithms suggests a focus on the learning or training phase before deployment, rather than inference. An example is the offline training of a classifier to discriminate different locomotion tasks using a prerecorded dataset from multiple subjects.

3.4.3. AI techniques: machine-learning approaches and logic- and knowledge-based approaches

Infer how to generate output The guidelines also emphasize the use of “how” in the expression “infer [...] *how* to generate output”, specifying that it goes beyond “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” which should be interpreted as referring to the “building phase, whereby a system derives outputs through AI techniques enabling inferencing” [3, p. 5]. In other words, a key characteristic of an AI system is that it does not merely generate outputs, but it also “understands” how to generate outputs. However, this only occurs in the pre-deployment phase.

Machine learning and logic- and knowledge-based approaches Recital 12 mentioned two categories of AI techniques that are employed in the building phase:

1. “machine learning approaches that learn from data how to achieve certain objectives, and
2. logic- and knowledge-based approaches that infer from encoded knowledge or symbolic representation of the task to be solved.”

Machine learning approaches encompass “a large variety of approaches enabling a system to ‘learn’, such as supervised learning, unsupervised learning, self-supervised learning and reinforcement learning” [3, p.6]. The focus of these approaches is on the system’s ability to learn which is arguably in contrast with logic- and knowledge-based approaches that “[i]nstead of learning from data, [...] learn from knowledge including rules, facts and relationships encoded by human experts” [3, p.7]. Based on such knowledge, AI systems reason rather than learn “via deductive or inductive engines or using operations such as sorting, searching, matching, chaining” and include “knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning, expert systems and search and optimisation methods” [3, p.7].

Learning and reasoning The definition of AI techniques that infer outputs from inputs seems thus very inclusive of a variety of approaches that either learn or reason. This distinction echoes a key difference between data and knowledge: data refer to raw measurements or observations (e.g., sensor readings), whereas knowledge consists of structured, interpretable information derived from or imposed on data, often in the form of symbolic representations or semantic rules. Rather than learning through exposure to data, these systems operate through reasoning, applying predefined knowledge to novel situations [10]. For instance, in the control of powered prostheses, a knowledge-based system might rely on explicit rules such as “if this, then that” encoded by an expert. These systems interpret data acquired by the data sensor according to pre-encoded logic, without generalizing from new data. In contrast, a machine learning system learns to recognize gait phases or classify locomotion tasks by analyzing large volumes of recorded data. Although this creates a clear methodological distinction, the boundary is increasingly blurred in practice: emerging hybrid systems combine symbolic reasoning with data-driven learning (e.g., using ML to learn gait patterns while applying rule-based logic to ensure safety or interpretability), leveraging the strengths of both paradigms in assistive and rehabilitation technologies [11].

3.4.4. Systems that fall outside the scope of AI system definition

Owing to the explicit mention in Recital 12 that the AI system definition does not cover “simpler traditional software systems or programming approaches and [...] systems that are based on the rules defined solely by natural persons to automatically execute operations”, the guidelines mention four exceptions of systems that “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” (p. 8). The distinctive ability to adjust the output autonomously appears to refer to the self-learning capacity outlined in Section 3.3. This capacity arguably relates more to machine learning approaches than to logic- and knowledge-based ones.

No. 1: optimization. The first exception concerns those systems that are used to “improve mathematical optimisation or to accelerate and approximate traditional, well-established optimisation methods, such as linear or logistic regression methods” [3, p.8]. Optimization-based systems determine the best-fitting relationship between input variables and outcomes by optimizing a defined objective function. Although these methods are often associated with machine learning, they are rooted in classical statistical modeling and are based on deterministic optimization rather than on learning from data[12]. In our view, this description does not raise any interpretative issue: when systems only perform optimization, they are excluded from the AI Act.

No. 2: basic data processing. The second case refers to “basic data processing systems” that follow “predefined, explicit instructions or operations” [3, p.9]. This means systems that “execute tasks based on manual inputs or rules, without any ‘learning, reasoning or modelling’ at any stage of the system lifecycle” (p. 9). The AI Office further clarifies that this category of systems operates “based on fixed human-programmed rules, without using AI techniques [...] to generate outputs” (p. 9), which appears to be tautological, since it states that the definition of AI systems excludes those systems that do not use

AI techniques. The expression “basic data processing systems” is not commonly employed in computer science and engineering practice. Still, it likely refers to operations of importing, exporting, extracting, filtering and reviewing data [13]. The examples provided in the guidelines, indeed, refer to database management systems that sort or filter data (e.g., MongoDB, MySQL, etc), standard spreadsheet software applications (E.g., Excel), software that performs descriptive analysis, hypothesis testing (e.g., R, Stata, etc), and visualization. Even though the notion of “basic data processing systems” is broad and not standardized, we believe that the examples reasonably clarify the scope of application of this exception.

No. 3: classical heuristics. According to the guidelines, the definition further excludes systems based on classical heuristics, characterized as “problem-solving techniques that rely on experience-based methods to find approximate solutions efficiently” (p. 9). Classical heuristics can be seen as a subset of inference techniques. Whereas the latter encompasses a broad range of reasoning methods, including those that evolve or adapt over time, classical heuristics are typically more rigid and do not adapt or learn from new data (as indeed mentioned in the guidelines), unlike modern inference techniques [14]. Therefore, classical heuristics can be considered a limited, rule-based form of inference and “typically involve rule-based approaches, pattern recognition, or trial-and-error strategies rather than data-driven learning” [3, p.9]. There seems to be an ontological discrepancy between these examples, though, because pattern recognition (which uses statistical methods to discern regularities and structures in the data, identifying recurring themes or trends) can be part of classical heuristics, but can also be a data-driven process, especially in modern machine learning systems. In contrast, classical heuristics may employ simpler, experience-based methods to recognize patterns without relying on data learning [15]. Thus, including pattern recognition as a typical example of classical heuristics may be confusing.

No. 4: simple prediction systems. Finally, “simple prediction 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 [their] performance” [3, p.10]. This definition is quite cryptical, since the meaning of “performance”, which is essential in the distinction, is not illustrated. The term could refer to either the outcome (e.g., accuracy) or the behavior of the system. However, no benchmark or threshold is provided to support this distinction. For example, in the case of the high-level control layer of robotic prostheses, which classifies locomotion tasks, “performance” may be interpreted as classification accuracy. However, without a reference point, this criterion is potentially arbitrary and lacks objectivity.

Lastly, it is possible to recover a standard definition of “statistical learning” (i.e., a set of methods and tools designed to understand complex data by extracting patterns and making predictions [16]) but not of “basic statistical learning”. Similarly, the notion of “simple prediction systems”, as opposed to, presumably, complex prediction systems is not straightforward either.

4. Discussion

4.1. Terminological and interpretative issues arising from the interdisciplinary analysis

Even though certain constitutive elements of the definition of AI system, as expressed in Article 3(1) and Recital 12, are clarified in the Guidelines (e.g., the exclusion of rule-based systems, of optimization techniques, etc), several issues of semantics and interpretation have emerged from our analysis (summarized in Table 1 in the Appendix):

1. The definition of autonomy that refers to independence of action looks rather vague and does not provide for the different components of an AI system that may have varying levels of autonomy, also depending on whether the perspective of the developer or the user is favored

2. whereas the guidelines suggest a complete overlap between the two categories, self-learning systems are only a subset of adaptive systems (plus, this ability pertains to the building phase rather than the use phase)
3. Even though the guidelines mention a clear-cut distinction, the difference between machine learning and logic- and knowledge-based approaches is often blurred in practice
4. The use of terms such as ‘learning’ or ‘training’ would be more appropriate than ‘inference’ in the machine learning domain
5. Pattern recognition is often based on data-driven learning; thus, it is confusing to mention it as a typical example of rule-based system as an alternative to it
6. The term “performance” of the system is not defined, even though it is key to understanding the scope of the exemption, and may refer either to the outcome or the behavior of the system
7. Several terms (“simple(r)”, “basic”, “traditional”) that are confidently used in the guidelines do not have a standard meaning in the engineering and computer science communities. Thus, it may be challenging to understand what they refer to.

As in other interdisciplinary regulatory areas (e.g., data protection), there is a need to agree on a consistent set of terms that describe AI systems, AI techniques, their legal and technical requirements, and so on. Such an agreed-upon vocabulary should become a shared language among various domain experts, facilitating communication and discussion of knowledge across disciplinary boundaries, regardless of the specific context in which AI systems are developed or used. We therefore call for clarification of the terms employed in the guidelines. As recalled in Recital 12, without a reliable vocabulary that can be used unequivocally and safely across communities, there may be risks to legal certainty that could hamper scientific research and technological development. The comparative vocabulary presented in these pages aims to constitute a first step towards creating a shared, interdisciplinary vocabulary that reliably drives compliance and achieves harmonization.

4.2. From definition to exceptions or from exceptions to definition?

From an operational point of view, our team has discussed at length whether it is more meaningful to determine if the prosthesis falls under the definition of AI system (i.e., the general rule) and then understand if it coincides with one of the exceptions listed in Section 3.4.4 or, rather, if we should immediately determine whether it falls under one of the exceptions, which would have the benefit of rapidly ruling out the applicability of the AI Act.

We lean towards the first method because the guidelines define what is *not considered* an AI system under the AI Act. This is conceptually different from affirming that a machine-based system *is not* an AI system. We believe that there is a high probability that an automated system may be considered an AI system in the meaning of the AI Act, since the seven conditions do not need to be present in both the pre-deployment and post-deployment phases. Further, from a logical point of view, exceptions are elements that belong to a certain category, but due to specific (policy) reasons, they are not included in the general rule. Regardless of whether a system is considered an AI system in the technical community, the policy definition is the relevant factor in this analysis.

4.3. Relevance of this work

As soft law instruments, the guidelines of the AI office should inform compliance activities, by providing reliable criteria that various stakeholders can use to evaluate AI systems and take appropriate actions to address the requirements (e.g., in the case of developers and deployers) or to oversee if these have been correctly implemented (e.g., in the case of supervisory authorities). Guidelines should thus support a responsible and accountable approach to technological innovation. However, we have shown that finding a coherent interdisciplinary method for the interpretation of the connections between the AI Act, the guidelines, technical terms, and concrete use-cases is challenging. In interdisciplinary domains, to close the gap between policy terms and technical definitions, the method of interpretation should

rely on a common vocabulary, combined with an attentive analysis of the grammar and the overall logical structure of the text.

With the answers provided here (preliminary) on what falls under the scope of the AI Act and what does not, we aim to offer guidance to other researchers and practitioners who share similar doubts, regardless of their use cases. Developers and engineers may find it valuable to analyze (and, if necessary, complete or review) the definitions we have formulated and mapped to the terminology employed in the legal sources. Legal scholars, policy-makers, and regulators may find it meaningful to build on our interpretations to determine and solve the issues of applicability of the AI Act to real-world machine-based systems. This is paramount due to the consequences of such a classification. If a software is understood as an AI system and is used as a medical device, it will be classified as a high-risk system, as set forth by Article 6(1) and Annex I (11). This means that such AI systems will not only need to comply with the Medical Devices Regulation (EU Regulation 2017/745), but also with the many requirements of the AI Act. However, it is currently unknown how to coordinate compliance efforts towards both regulations without overburdening developers.

This interoperable interpretation framework for robotic prostheses should be further improved with different, more complex use cases to achieve generalization. Our efforts are meant not only to address the ethical-legal compliance in a given R&D life-cycle, but also to contribute to pre-standardize and develop interoperable tools of interpretation for a fragmented, ever-evolving legal framework that governs cutting-edge technological development. We maintain that this work is a necessary precondition to an accountable, future-proof approach.

5. Conclusions and future work

This article describes early work of an ongoing interdisciplinary dialogue between bioengineering researchers and legal scholars. The commented, comparative lexicon we have elaborated may evolve as our collaboration continues, since we need to determine whether the robotic prosthesis described in Section 2 falls within the definition of an AI system. If it does, the robotic prosthesis would be classified as a high-risk AI system because it is a medical device. Therefore, even though it is developed within research settings, it will need to incorporate several legal and technical requirements early on, which will later enable it to be put on the market. We welcome constructive feedback from the readers of this article and our colleagues at the HHAI's Workshop on Law, Society, and Artificial Intelligence to build a reliable knowledge base of AI that spans across domains.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT (developed by OpenAI) for sentence polishing.

Acknowledgments

This work was supported by the Italian Ministry of Research (Biorobotics Research and Innovation Engineering Facilities (BRIEF), GA IR0000036 – CUP J13C22000400007 and Fit4MedRob - Fit for Medical Robotics, GA PNC0000007); the European Commission (SoBigData PPP RI Preparatory Phase Project, GA 101079043 and Smart Maritime and Underwater Guardian (SMAUG), GA 101121129); and the Istituto Nazionale Assicurazioni Infortuni sul Lavoro (INAIL) (MOTU++, GA PR19-PAI-P2). We thank the anonymous reviewers for their helpful comments.

References

- [1] C. Sganga, D. Amram, F. Gennari, A. Rossi, A. Blatti, P. Turan, D.7.4 cross-field regulatory analysis (phase ii) (2024). URL: <https://zenodo.org/records/14655644>.
- [2] I. Fagioli, A. Mazzarini, C. Livolsi, E. Gruppioni, N. Vitiello, S. Crea, E. Trigili, Advancements and challenges in the development of robotic lower limb prostheses: a systematic review, *IEEE Transactions on Medical Robotics and Bionics* (2024).
- [3] E. Commission, Annex to the communication to the commission approval of the content of the draft communication from the commission guidelines on the definition of an artificial intelligence system established by regulation (eu) 2024/1689 (ai act), 2025. URL: <https://ec.europa.eu/newsroom/dae/redirection/document/112455>.
- [4] I. Fagioli, A. Mazzarini, A. Baldoni, F. Dell'Agnello, E. Gruppioni, E. Trigili, S. Crea, N. Vitiello, A lightweight robotic knee prosthesis with torsional series elastic actuation, in: 2024 10th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob), IEEE, 2024, pp. 389–394.
- [5] A. Mazzarini, I. Fagioli, A. Baldoni, F. Dell'Agnello, E. Gruppioni, E. Trigili, S. Crea, N. Vitiello, A robotic ankle-foot prosthesis based on torsional series and parallel elasticity, in: 2024 10th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob), IEEE, 2024, pp. 395–400.
- [6] M. R. Tucker, J. Olivier, A. Pagel, H. Bleuler, M. Bouri, O. Lamercy, J. d. R. Millán, R. Riener, H. Vallery, R. Gassert, Control strategies for active lower extremity prosthetics and orthotics: a review, *Journal of neuroengineering and rehabilitation* 12 (2015) 1–30.
- [7] I. Mareels, J. W. Polderman, *Adaptive Systems*, Birkhäuser Boston, Boston, MA, 1996, pp. 1–26. doi:10.1007/978-0-8176-8142-5_1.
- [8] A. G. B. Richard S. Sutton, *Reinforcement Learning: An Introduction*, 2nd Edition, 2 ed., Bradford Books, 2018.
- [9] ISO, *Iso/iec 22989:2022 information technology — artificial intelligence — artificial intelligence concepts and terminology*, 2022. URL: <https://www.iso.org/standard/74296.html>.
- [10] S. J. Russell, P. Norvig, *Artificial intelligence: a modern approach*, pearson, 2016.
- [11] R. Gehlhar, M. Tucker, A. J. Young, A. D. Ames, A review of current state-of-the-art control methods for lower-limb powered prostheses, *Annual Reviews in Control* 55 (2023) 142–164.
- [12] S. P. Boyd, L. Vandenberghe, *Convex optimization*, Cambridge university press, 2004.
- [13] Y. Xia, J. Sun, *Basic Data Processing in QIIME 2*, Springer International Publishing, Cham, 2023, p. 65–94. URL: https://doi.org/10.1007/978-3-031-21391-5_3. doi:10.1007/978-3-031-21391-5_3.
- [14] J. Pearl, *Heuristics: intelligent search strategies for computer problem solving*, Addison-Wesley Longman Publishing Co., Inc., 1984.
- [15] S. Pal, A. Pal, *Pattern recognition. from classical to modern approaches*, World Scientific Pub Co Inc (Printed in Singapore) (2001).
- [16] G. James, D. Witten, T. Hastie, R. Tibshirani, et al., *An introduction to statistical learning*, volume 112, Springer, 2013.

A. Summary of critical interpretative issues

Art. 3(1)	Recital 12	Guidelines	Issue
“Varying levels of autonomy”	“some degree of independence of action”	Ability to generate an output “on its own”	Each component may have its own level of autonomy
“May exhibit adaptiveness”	adaptiveness “refers to self-learning capabilities”	“a system may possess adaptiveness, but not necessarily, <i>or</i> self-learning capabilities after deployment”	Not all adaptive systems have self-learning abilities
“infer [...] <i>how</i> to generate output”	“The techniques that enable inference while building an AI system include machine learning [...] and logic- and knowledge-based approaches”	ML approaches encompass “a large variety of approaches enabling a system to ‘learn’ ” whereas logic- and knowledge-based approaches “[i]nstead of learning from data, [...] learn from knowledge including rules, facts and relationships encoded by human experts”	Increasingly blurred distinction between the two approaches
“infer [...] <i>how</i> to generate output”	“This capability to infer refers to [...] a capability of AI systems to derive models or algorithms, or both, from inputs or data”	which “underlines the relevance of the techniques used for building a system”	In ML it would be more appropriate to use terms such as learning or training, rather than inference.
[AI systems that do not infer (exception no. 3)]	“simpler traditional software systems or programming approaches and [...] systems that are based on the rules defined solely by natural persons”	e.g., classical heuristics “typically involve rule-based approaches, pattern recognition, or trial-and-error strategies rather than data-driven learning”	Pattern recognition may also be data-driven
[AI systems that do not infer (exception no. 4)]	see above	“simple prediction systems [...] whose performance can be achieved via a basic statistical learning rule, [...] fall outside the scope of the AI system definition, due to [their] performance”	Performance is not defined and may indicate either the results or the behavior of the system
-	-	“simpler traditional software systems or programming approaches”, “basic data processing systems”, “simple prediction systems”, “reasonable degree”	These terms are not commonly used within the engineering / computer science community

Table 1

Table summarizing the terminological and interpretative issues arising from the combined analysis of article 3(1), Recital 12, Guidelines [3] and technical definitions