

Accepted version.

The publisher's version can be found at <https://doi.org/10.1109/CAI59869.2024.00179>

“© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

# Navigating the EU AI Act: A Methodological Approach to Compliance for Safety-critical Products

Jessica Kelly, Shanza Ali Zafar, Lena Heidemann, João-Vitor Zacchi, Delfina Espinoza, Núria Mata

Fraunhofer IKS

Munich, Germany

jessica.kelly@iks.fraunhofer.de

**Abstract**—In December 2023, the European Parliament provisionally agreed on the EU AI Act. This unprecedented regulatory framework for AI systems lays out guidelines to ensure the safety, legality, and trustworthiness of AI products. This paper presents a methodology for interpreting the EU AI Act requirements for high-risk AI systems by leveraging product quality models. Our methodology facilitates the development and assessment of Artificial Intelligence (AI) systems that not only adhere to established quality standards, but also comply with the regulatory requirements outlined in the Act for high-risk (including safety-critical) AI systems. We demonstrate the applicability of this methodology on an exemplary automotive supply chain use case, where several stakeholders interact to achieve EU AI Act compliance.

**Index Terms**—EU AI Act, Quality Attributes, AI Systems, Regulations.

## I. INTRODUCTION

With the growing number of AI systems being deployed in safety-critical applications, there is a pressing need to establish regulations that govern the safe and responsible use of AI. In December 2023, the EU AI Act (Act) was provisionally agreed upon by the European Parliament, setting the precedent for the regulation of AI applications [1]. It is the first comprehensive regulatory framework governing the development, deployment, and use of AI systems. The Act introduces a risk-based classification of AI products. Applications whose risk is deemed *unacceptable*, such as social-scoring systems, are banned within the framework of the Act. Applications with a risk rated *high* (high-risk) must demonstrate compliance with stringent requirements ensuring that, among others, safety, transparency, and human rights needs are met. The outlined requirements affect not only AI products, but any stakeholders involved in the AI value chain. Organizations will need to adapt to the evolving landscape of the Act, balancing innovation and regulatory adherence. While entities across the AI value chain will need to align with the framework of the Act, future standards and regulations will also be affected.

Safety-critical systems, whose failure could result in significant harm to people or even loss of life, are deemed high-risk under the Act. Existing domain-specific safety standards, such as ISO 26262 [2] for automotive or ARP 4761 [3], DO-178C

[4], and DO-254 [5] for aerospace, cover some aspects of the Act requirements for high-risk AI systems. However, these standards do not, in their current state, address AI specific considerations for safety, transparency, and human oversight. Although efforts are being made to develop new safety standards for AI systems, their development and adherence are a time-intensive process. Quality models for AI products, such as ISO/IEC 25059:2023 [6], can help address the requirements set out in the Act early in the development cycle. Additionally, they provide the flexibility to include attributes that may not be safety relevant, but ensure better quality.

The Act introduces additional challenges to compliance when several stakeholders are involved. Safety-critical AI products are typically part of complex global supply chains, where many suppliers interact to produce the final product. In the automotive industry, for example, it is rare that a single entity is responsible for the development and integration of all vehicle components. Demonstrating compliance to the Act thus becomes an increasingly challenging and intricate task. It is clear that organizations will need tools and methodologies to address the requirements laid out by the regulation. Specifically, a systematic methodology that aids organizations in verifying compliance is required. To facilitate this, our work leverages product quality models to break down the Act requirements into verifiable properties. In the first phase, an extended quality model for AI systems is derived using attributes relevant to the Act. Next, using this quality model, an approach to map the articles of the Act to quality attributes for AI Systems is presented. Finally, to address the complexities arising from supply chain relationships, a contract-based approach for the derivation of technical requirements from quality attributes is proposed. This methodology is, to the best of our knowledge, the first systematic approach for deriving technical requirements at the stakeholder level from high-level Act requirements.

The contributions are as follows:

- An extended quality model for safety-critical AI systems, which covers relevant attributes for the EU AI Act;
- A systematic approach for mapping the Act requirements to relevant quality attributes in the extended quality model;
- A contract-based approach for deriving verifiable technical requirements for the quality attributes; and finally,
- An exemplary use case for an automotive supply chain is

This work was funded by the Bavarian Ministry for Economic Affairs, Regional Development and Energy as part of a project to support the thematic development of the Institute for Cognitive Systems.

presented to demonstrate the applicability of the proposed methodology.

This paper is organized as follows: Section II provides an overview of related work. Section III outlines the details of our proposed methodology for addressing the Act requirements for High-Risk AI systems. In Section IV, we apply our methodology to an Automotive Supply Chain Use-Case. We discuss the implications of our methodology in Section V and Section VI summarizes our contributions and future work.

## II. BACKGROUND

Given the novelty of the Act, there is currently little work surrounding how the requirements for high-risk AI systems should be addressed. Many organizations seek to understand whether compliance with current regulations can assist in addressing the EU AI Act. Existing standards do not fully cover the stringent requirements laid out in the Act, such as transparency, lawfulness, and fairness. Product quality models may help fill this gap, and can be adapted to include properties that may not be safety relevant but do contribute to quality. Existing quality standards, such as ISO/IEC 25010 Product Quality Standard [7] and ISO/IEC FCD 25012 Data Quality Model [8] do not address AI specific attributes such as transparency, controllability, and intervenability. The Quality Model for AI Products/Systems in ISO 25059:2023 [6] introduces some AI specific attributes like functional adaptability, and robustness, however, it is still lacking in its coverage of attributes like transparency, monitorability, and data quality, among others. ISO/IEC 24028 - *overview of trustworthiness in artificial intelligence* highlights the need for new standards which incorporate AI specific quality attributes [9]. Given this, recent contributions have addressed the need for extended quality models for AI systems.

The authors of [10] define a systematic process for deriving a quality model for ML systems. They formalize the derivation of quality attributes using a quality meta-model, enabling the modelling of different hierarchies of quality. From this meta-model, relevant entities are defined and categorized into corresponding views of an ML product, namely the model, data, infrastructure, environment, and system views. Relevant properties are then described for a selected use case, and a list of corresponding metrics is proposed. This systematic approach ensures a comprehensive coverage of ML-related quality properties, however, it may not be well suited for addressing the Act. Given the high-level nature of the Act, it is beneficial to address high-level properties of AI products, which may apply to several levels of abstraction and stakeholder perspectives. In addition, alignment with existing standards is relevant for organizations wishing to understand their current coverage in their development practices. As such, the extended quality model proposed in this paper is based on an alignment to high-level product quality standards, and other existing safety standards.

Aside from quality models, recent literature has emerged proposing different approaches for addressing the Act. Novelli et al. [11] highlight the importance of accurately assessing the

risk of AI systems in the context of the Act. The authors introduce a risk-assessment model to improve the accuracy of this risk estimation for ethical and safe AI practices in accordance with the Act. While relevant to addressing the Act, this work focuses only on the risk classification and does not delve into the requirements for AI systems that are deemed high-risk. A different perspective is taken in [12], which provides an overview of explainability requirements in the Act, proposing metrics for assessing AI Act compliance. The authors highlight the need for metrics that are risk-focused, model-agnostic, goal-aware, intelligible, and accessible, and assess current metrics against these criteria. The paper provides a thorough coverage of explainability, but does not address the broader spectrum of requirements outlined in the EU AI Act. It does not address the need for a comprehensive methodology for addressing these requirements from the perspective of different stakeholders, leaving a gap in practical guidance for entities seeking compliance. A more pragmatic approach to compliance is suggested in [13], where the authors propose a methodology for organizations to measure their compliance to the Act using a comprehensive questionnaire. However, the approach focuses on measuring compliance to the Act, and does not provide guidance to organizations who may seek further compliance.

## III. PROPOSED METHODOLOGY: ELICITING HIGH-LEVEL REQUIREMENTS FROM THE ACT

This section presents the systematic methodology for eliciting high-level requirements from the EU AI Act. First, an overview of the extended quality model for AI products is presented, followed by the approach for mapping Act requirements to quality attributes. Finally, a contract-based approach for deriving technical requirements for quality attributes is proposed.

### A. Deriving an Extended Quality Model for Safety-critical AI Systems

To derive relevant quality attributes for safety-critical AI systems, ISO/IEC 25059 [6] is used as a baseline. ISO/IEC 25059 provides the quality model serving as an extension to the ISO 25010:2011 series - Systems and Software Quality Requirements and Evaluation (SQuaRE) [7]. ISO/IEC 25059 establishes consistent terminology for specifying, measuring, and evaluating the quality of AI systems. It considers the quality model from two perspectives, product quality and quality in use. In this report, we will only focus on the product quality model. The product quality model from ISO/IEC 25059 is highlighted in Fig. 1.

We extend the product quality model presented in ISO/IEC 25059, with a specific focus on the following points:

- Covering relevant topics from the Act to increase trustworthiness. ISO/IEC 25059 has some gaps when it comes to the coverage of Act requirements, for example, there is a lack of consideration for human oversight, transparency for different stakeholders, and ethical integrity. We have added them as attributes in the extended quality model.

- Integrating safety and data quality attributes in the ISO/IEC 25059 product quality model. The safety attribute, present in ISO/IEC 25010:2011 upon which ISO/IEC 25059 is based, is notably absent in ISO/IEC 25059. Similarly, the data quality model is extended from ISO/IEC 25012:2008 [8]. We have included it in our extended quality model due to the high impact data quality has on the quality (including safety) of AI systems.
- Incorporating AI-related safety properties and data quality from other sources, such as work from [14], or the upcoming safety standard for AI systems in road vehicles, ISO PAS 8800 [15].
- Aligning ISO/IEC 25059:2023 with the updated version of ISO/IEC 25010:2023. It is currently based on ISO/IEC 25010:2011.

The extended model is depicted in Fig. 1. Definitions for quality attributes and sub-attributes are given in Table I. This methodology can be adapted as new standards emerge regarding AI system product quality. For the safety attribute, we recommend using domain-specific standards, where available, for a more systematic approach. For instance, combining ISO 26262 [2], ISO 21448 [16], and the upcoming ISO PAS 8800 [15] for road vehicles.

### B. Mapping EU AI Act Articles to the Extended Quality Model

The Act articles for high-risk AI systems do not provide guidelines for achieving compliance. We propose to map these articles to our extended quality model. Such a mapping can be leveraged to assess the coverage of the Act based on measurable properties of AI systems. We used our own experiences and research, coming from diverse research backgrounds, to consolidate a detailed mapping of high-level requirements to quality attributes. A high-level summary of the mapping is shown in Table II.

### C. Contract-Based Validation Approach for Quality Attributes

High-risk AI applications are typically part of complex global supply chains, in which several stakeholders are involved. In this context, ensuring the fair, lawful, and ethical development of AI applications is notably challenging. Parallels can be drawn with the recently enacted Supply Chain Act for companies headquartered in Germany [21]. This legislation extends the responsibility of organizations and mandates the safeguarding of human rights and environmental protection throughout the entire supply chain. While not specific to AI, this legislation provides insights into how a company's responsibility for regulatory adherence is not simple, and may in some cases include indirect suppliers. A similar perspective can be applied to the Act where the responsibility is defined for some actors within the AI value chain, yet remains unspecified for others.

The Act defines a set of relevant AI actors, and outlines responsibilities for compliance depending on these defined roles. Principal responsibility for compliance is assigned to the provider of a high-risk AI system. However, in the case of safety-critical systems, any manufacturer in the supply chain

TABLE I  
NEW OR MODIFIED DEFINITIONS OF QUALITY (SUB-)ATTRIBUTES IN THE EXTENDED QUALITY MODEL (SEE FIG. 1).

| Term               | Definition  |
|--------------------|---|
| Ethical Integrity  | The extent to which an entity's actions, beliefs, methods, measures, and principles all derive from a single core group of values.  |
| Human Oversight    | The ability for humans to understand, supervise, and control the design and operation of AI-based systems. [17]   |
| Fairness           | The extent to which a system prevents unjust predictions towards protected attributes (race, gender, income, etc). Ability of the model to output fair decisions. [10]  |
| Privacy Protection | The extent to which the product or system protects the privacy and handles sensitive information of the stakeholders involved (users, people in training examples).   |
| Value Alignment    | The extent to which the AI system behaviour is aligned with human values. [15]  |
| Self-Monitoring    | The extent to which the system is aware of its state so it can respond appropriately to avoid going to a harmful state.   |
| Documentability    | see ISO/IEC/IEEE 24765 [18]   |
| User Transparency  | Degree to which the functionalities of the system are clear to the intended user.   |
| Interpretability   | The extent to which the inner workings of the AI system can be analyzed in order to understand why it behaves the way it does.  |
| Traceability       | The extent to which there exists data and processes that can record the system's decisions and link artifacts at different stages. [19]   |
| Explainability     | see ISO 22989 [20]  |
| Accountability     | Capability of a product to enable actions of a human to be traced uniquely to the human.  |
| Monitorability     | The extent to which relevant indicators of an AI system are effectively observed/monitored and integrated in the operation of the system.   |
| Representativeness | The distribution of data (or probability of distribution) truly corresponds to the information in the environment or the phenomenon to be captured. [15]  |
| Independence       | The data at a specific level of architectural abstraction are not affected by changes to lower levels of abstraction. Separate sets of data are used for specific purposes where required (e.g. AI training data, AI validation data). [15] |
| Data Fairness      | Degree to which the data is free from bias against a given group. [10]  |
| Availability       | The degree to which data has attributes that enable it to be retrieved by authorized users and/or applications in a specific context of use and within the time required. (see [8] and [15])  |
| Integrity          | The data are unaltered either by natural phenomenon (e.g. noise) or intentional action (e.g. poisoning). [15]   |
| Temporality        | A general property referring to temporal characteristics of data e.g. its timeliness, ageing or lifetime. [15]  |

can also be assigned responsibility. Importers and distributors are required to verify that a provider has met their obligations. Given the complexities arising from an ambiguous assignment of responsibilities, stakeholders will likely need to ensure not only their own compliance, but in certain cases the compliance of other involved actors.

One of the few approaches to deriving a use-case agnostic, stakeholder-specific approach to compliance is provided in the EU Model contractual clauses [22] for the procurement of AI systems from external stakeholders. The clauses are generic and adaptable to specific use cases, and provide

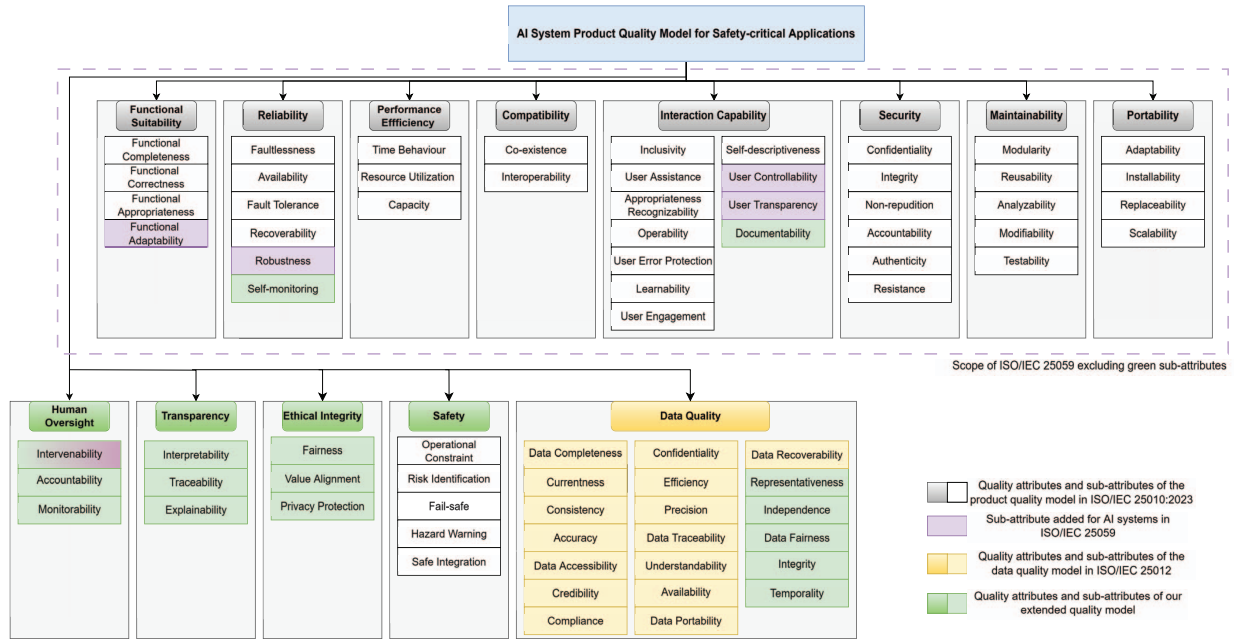


Fig. 1. Extended quality model for AI products for safety-critical applications.

organizations wishing to procure AI systems with a possible solution to ensuring compliance with the Act. The clauses are aligned with the Act, and support the ethical, transparent, and accountable development of AI. The Commission highlights that these clauses may need to be adjusted depending on the contractual relationships, as they do not distinguish between the obligations of the many actors discussed in the Act. Additionally, there is a need for a concrete methodology to derive technical requirements from such contractual clauses. We propose a contract-based approach for the systematic validation of the Act requirements across the value chain.

Our approach is based on design contracts. Design contracts define *guarantees* which are guaranteed to be fulfilled by the stakeholder. The fulfillment of said design contract is only guaranteed given that a set of *assumptions* is fulfilled [23]. Verifying EU AI Act compliance thus boils down to the interface with the design contracts. Given that all *assumptions* are fulfilled, *guarantees* are assumed to be fulfilled. We demonstrate this approach using an exemplary automotive supply chain use case, shown in Fig. 2. For the sake of simplicity, we consider a car manufacturer which integrates (n) sub-systems. Each stakeholder in the supply chain may come from different entities. The design contracts are shown for each relevant stakeholder. The technical requirements are derived from the *assumptions* and flow between stakeholders. Stakeholder definitions are taken from [24] and [20]. An example of this validation approach for a chosen quality attribute is presented in the following section.

TABLE II  
MAPPING OF EU AI ACT REQUIREMENTS TO QUALITY ATTRIBUTES FOR SAFETY-CRITICAL AI SYSTEMS.

| Article   | Sub-Attribute Mapping   |
|---|---|
| <b>9. Risk management system</b>                              | Risk identification, Testability, Value Alignment   |
| <b>10. Data and data governance</b>                           | Independence, Data Completeness, Currentness, Independence, Data Fairness, Precision, Representativeness, Consistency, Accuracy, Credibility, Temporality, Confidentiality, Compliance, Data Traceability |
| <b>11. Technical documentation</b>                            | Traceability  |
| <b>12. Record-keeping</b>                                     | Operability, Non-repudiation, Traceability, Self-descriptiveness, Accountability, Self-Monitoring, User Engagement, Monitorability  |
| <b>13. Transparency and provision of information to users</b> | User Engagement, Self-descriptiveness, User Transparency, Interpretability, Documentability, Appropriateness Recognizability  |
| <b>14. Human oversight</b>                                    | Documentability, Learnability, Value Alignment, Accountability, Interpretability, Fairness, Explainability, Intervenable, Monitorability, User Error Protection.  |
| <b>15. Accuracy, robustness, and cybersecurity</b>            | Functional Correctness, Faultlessness, Robustness, Appropriateness Recognizability, Self-descriptiveness, Functional Adaptability, Fault Tolerance, Robustness, Integrity, Resistance                     |

#### IV. USE-CASE DEMONSTRATION: AUTOMOTIVE SUPPLY CHAIN

To demonstrate the applicability of our contract-based validation approach, we consider the typical automotive supply




Fig. 2. Automotive supply chain demonstrating stakeholder interactions and respective design contracts (DCs) and technical requirements (TRs).

chain presented in Fig. 2. Suppose we have a Traffic Sign Recognition (TSR) component as a sub-system for a car manufacturer, as depicted in Fig. 3. We would like to verify compliance for a given quality attribute in Table II. Due to its applicability to *Article 14: Transparency and Provision of Information to Users* and *Article 15: Human Oversight*, we select Explainability (for definition see Table I) as our quality sub-attribute. Starting with the AI Product or Service Provider as our primary stakeholder, we would see stakeholder interactions as illustrated in Fig. 3.




Fig. 3. Design contracts and technical requirements elicited by the AI product or service provider.

We first specify the design contract for the primary stakeholder. We define the *assumptions* that are required so that the primary stakeholder can demonstrate compliance to the Act. In this case, we define *assumptions* for explainability of the AI component. These *assumptions* are exemplary and would be refined depending on the relevant use case. From these *assumptions*, the AI product or service provider would define *guarantees* that it can satisfy, given that the *assumptions* are met. An example of a design contract for the AI Product or Service Provider is shown in Table III.

The AI Provider's *assumptions* would be detailed as technical requirements for the relevant stakeholders. In Table IV, we provide examples of how these requirements might be formulated from the technical point of view in a legal contract. This approach provides a formal method to derive

TABLE III  
DESIGN CONTRACT FOR THE AI PRODUCT OR SERVICE PROVIDER.

| Assumptions |   |
|-------------|---|
| 1           | The TSR component can be analyzed to understand its behavior. Documentation with global class-wise explanations is provided and representative. |
| 2           | Appropriate documentation regarding the development of the TSR model is available.  |
| 3           | The TSR can express important factors influencing its predictions in a way that humans can understand.  |
| 4           | Documentation from the AI system integrator regarding how sub-systems interact in the overall car is available.                                 |
| Guarantees  |   |
| 1           | Appropriate documentation regarding the design, development, licensing, and usage restrictions of the TSR is available.                         |

technical requirements for Act requirements using contract-based design.

TABLE IV  
REQUIREMENTS GIVEN BY THE AI PRODUCT OR SERVICE PROVIDER.

| Technical Requirement | Description   | Owner                |
|-----------------------|---|----------------------|
| TR1                   | The model architecture is well documented so that an expert user can understand the inner workings of the TSR component.  | AI Developer         |
| TR2                   | An ex-ante explanation is available for the user of the AI system. For example, documentation containing global class-wise explanations is provided, using a state-of-the-art explainability method.            | AI Developer         |
| TR3                   | Documentation containing train/test/validation data, pre- and post-processing operations, optimization method, loss function, and hyperparameters used for training, is available.                              | AI Developer         |
| TR4                   | An ex-post explanation is available for the user of the AI system which satisfies the required level of explainability. For example, a local, post-modelling explainability method such as SHAP is implemented. | AI Developer         |
| TR5                   | The AI system integrator shall provide requirements for the TSR interface within the system.  | AI System Integrator |

## V. DISCUSSION

This work presents a systematic methodology that can be used to assess the Act requirements from the perspective of different stakeholders. The proposed approach does not claim complete coverage of the Act. Instead, the extended quality model and the mapping should be subject to iterative refinement. This continuous improvement can occur as new insights emerge, regulatory frameworks evolve, or additional AI-specific attributes are identified or modified.

The mapping does not provide a measure of the degree of coverage of each article. The goal of the mapping at this stage is to highlight the utility of quality models for addressing properties of AI models not addressed by current standards. Extensions to both our model and our methodology are expected in future work.

The quantification of quality attributes remains a challenge. Current models lack precise metrics for evaluating critical aspects such as fairness, transparency, and adaptability in AI systems. This lack of metrics is particularly problematic in the context of contractual agreements, where clear and quantifiable measures are essential.

The practical implementation of certain quality attributes, such as human oversight, raises questions about the applicability of these requirements in real-world scenarios. In fully autonomous vehicles, the concept of oversight is unclear, necessitating a rethinking of how such systems are evaluated and regulated.

Additionally, while certain attributes were adequately defined for conventional software, their application to AI systems reveals new complexity. *Faultlessness* in AI, for instance, must consider the probabilistic nature of AI decisions, necessitating a redefinition that accounts for AI-specific error types and learning biases. This reassessment is crucial for ensuring that the extended model not only introduces new attributes for AI but also appropriately reinterprets existing ones to align with the unique characteristics and demands of AI technologies.

## VI. CONCLUSION

The EU AI Act is a transformative legislation which reshapes the global landscape of fair and ethical AI development. In this paper, we present a systematic methodology for addressing the requirements for high-risk AI products introduced in the Act. We develop an extended quality model for AI systems, and propose to map these quality attributes to the Act requirements. To address compliance, a contract-based approach for defining technical requirements is presented, ensuring that stakeholders across complex supply chains adhere to the EU AI Act regulations. Our design contracts foster a flexible and structured approach to compliance. This methodology allows researchers and practitioners to bridge the gap between existing quality models and the regulatory demands of the Act. This facilitates the development and assessment of AI systems that not only adhere to established quality standards but also comply with the regulatory requirements outlined in the Act.

## ACKNOWLEDGMENT

We would like to express our gratitude to Luciana Tanzarella for her valuable input throughout the research process.

## REFERENCES

- [1] European Commission, "Proposal for a Regulation of the European Parliament and of The Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts," 2021.
- [2] "ISO 26262:2018 - Road vehicles - Functional safety, 2nd Edition," 2018.
- [3] SAE International, "ARP4761 - Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment," 1996.
- [4] RTCA, "DO-178C: Software Considerations in Airborne Systems and Equipment Certification," 2011.
- [5] —, "DO-254: Design Assurance Guidance for Airborne Electronic Hardware," 2000.
- [6] "ISO/IEC 25059:2022: Software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — Quality model for AI systems," 2022.
- [7] "ISO/IEC 25010: Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — Product quality model," 2023.
- [8] "ISO/IEC 25012: Software engineering — Software product Quality Requirements and Evaluation (SQuaRE) — Data quality model," 2008.
- [9] "ISO/IEC 24028: Information technology Artificial intelligence — Overview of trustworthiness in artificial intelligence," 2020.
- [10] J. Siebert, L. Joeckel, J. Heidrich, K. Nakamichi, K. Ohashi, I. Namba, R. Yamamoto, and M. Aoyama, "Towards Guidelines for Assessing Qualities of Machine Learning Systems," in *Quality of Information and Communications Technology*, M. Shepperd, F. Brito e Abreu, A. Rodrigues da Silva, and R. Pérez-Castillo, Eds. Cham: Springer International Publishing, 2020, vol. 1266, pp. 17–31.
- [11] C. Novelli, F. Casolari, A. Rotolo, M. Taddeo, and L. Floridi, "Taking AI risks seriously: a new assessment model for the AI act," *AI & SOCIETY*, pp. 1–5, 2023.
- [12] F. Sovrano, S. Sapienza, M. Palmirani, and F. Vitali, "Metrics, explainability and the european ai act proposal," *J*, vol. 5, no. 1, pp. 126–138, 2022.
- [13] J. Walters, D. Dey, D. Bhamik, and S. Horsman, "Complying with the EU AI act," 2023.
- [14] J. Siebert, L. Joeckel, J. Heidrich, A. Trendowicz, K. Nakamichi, K. Ohashi, I. Namba, R. Yamamoto, and M. Aoyama, "Construction of a quality model for machine learning systems," *Software Quality Journal*, vol. 30, no. 2, pp. 307–335, Jun. 2022.
- [15] "ISO/PAS 8800: Road Vehicles - Safety and artificial intelligence," Tech. Rep., in work.
- [16] "ISO/PAS 21448: Safety Of The Intended Functionality – SOTIF," 2022.
- [17] European Commission, "Ethics By Design and Ethics of Use Approaches for Artificial Intelligence," 2021.
- [18] "ISO/IEC/IEEE 24765: International Standard - Systems and software engineering - Vocabulary," 2017.
- [19] G. Li, B. Liu, and H. Zhang, "Quality Attributes of Trustworthy Artificial Intelligence in Normative Documents and Secondary Studies: A Preliminary Review," *Computer*, vol. 56, no. 4, pp. 28–37, Apr. 2023.
- [20] International Organization for Standardization, "ISO/IEC 22989: Information technology — Artificial intelligence — Artificial intelligence concepts and terminology," 2022.
- [21] Federal Ministry of Labour and Social Affairs, "Act on Corporate Due Diligence Obligations in Supply Chains," January 2023.
- [22] P. R. Jeroen Naves, "Proposal for standard contractual clauses for the procurement of artificial intelligence (ai) by public organisations," September 2023.
- [23] A. Benveniste, B. Caillaud, D. Nickovic, R. Passerone, J.-B. Ralet, P. Reinkemeier, A. Sangiovanni-Vincentelli, W. Damm, T. A. Henzinger, and K. G. Larsen, *Contracts for system design*. Hanover, Massachusetts: Now Publishers, 2018.
- [24] N. N. G. d. Andrade and A. Zarra, "Artificial Intelligence Act: A Policy Prototyping Experiment: Operationalizing the Requirements for AI Systems – Part I," Rochester, NY, Nov. 2022.