

RESEARCH ARTICLE



Assessing the AI Act from the Root to the Leaf: A Decision Tree Methodology

Marinella Quaranta^{1,2,3,*} , Alessandro Bucci⁴ and Ilaria Angela Amantea¹

¹Department of Computer Science, University of Turin, Italy

²Department of Legal Studies, University of Bologna, Italy

³Vrije Universiteit Brussel, Belgium

⁴Department of Computer Science, University of Pisa, Italy

Abstract: The European Union's Artificial Intelligence Act (AI Act) establishes a risk-based regulatory framework that imposes progressively stringent requirements on AI systems in accordance with their associated risk levels. Despite its comprehensive scope, the Act's complexity poses significant challenges for organizations seeking to evaluate AI system risks and determine the specific regulatory obligations applicable to them. To mitigate these difficulties, this study proposes an integrated and methodical approach designed to enhance the accessibility and practical application of the AI Act. The proposed framework employs a decision tree-based methodology to systematically classify AI systems and map their corresponding regulatory obligations, thereby constructing a normative pathway aligned with the Act's evolving interpretative landscape. Furthermore, the approach incorporates Defeasible Deontic Logic to model the intricate relationships between legal norms, exceptions, and conditional obligations. The resulting method enables precise classification of AI systems and delineates their requisite regulatory obligations based on input data provided by AI developers and manufacturers, thereby supporting more transparent and compliant AI governance practices.

Keywords: AI Act, compliance, Defeasible Deontic Logic, decision tree

1. Introduction

Recent years have witnessed a profound technological transformation shaped by the integration of artificial intelligence (AI) into nearly every sector of the economy and society. This evolution, underpinned by breakthroughs in computational power, advances in algorithmic design, and the abundance of large-scale datasets, has redefined the dynamics of innovation and governance. In Europe, the framework to establish a trustworthy, transparent, and accountable approach to AI use is the Artificial Intelligence Act (AI Act or AIA).¹

The AI Act, however, remains an evolving, complex, and at times ambiguous regulatory framework. Initially simpler in its draft stages, the Act has gained layers of complexity over time, culminating in its final form in 2024. This complexity has led

to confusion and, in some instances, contradictions in the interpretation of its provisions. Even legal experts face challenges in parsing the nuances of the Act, raising concerns about the clarity and coherence of certain key definitions and criteria [1–3]. These ambiguities are particularly problematic for providers, especially those developing high-risk AI systems, who need a clear understanding of the Act's requirements.

To mitigate these compliance challenges, particularly for providers who may not have access to specialized legal resources, we propose a structured decision tree-based (DT-based) framework. This framework is designed to make the risk classification process set out in the AI Act easier to understand, allowing AI system providers, including those without legal expertise, to classify AI systems more precisely and efficiently. Through this method, we seek to address the ambiguity in the Act's risk classification, particularly regarding multi-category risk scenarios where an AI system could feasibly be assigned to more than one risk level.

At its core, the DT-based framework is a tool that combines Defeasible Deontic Logic (DDL) and Clingo logic programming with decision tree (DT) methodology to enhance the accessibility, reliability, and reproducibility of the classification process. The integration of DDL enables the creation of rule-based criteria that reflect the structure of the AI Act, while Clingo offers

*Corresponding author: Marinella Quaranta, Department of Computer Science, University of Turin, Italy, Department of Legal Studies, University of Bologna, Italy, and Vrije Universiteit Brussel, Belgium. Email: marinella.quaranta@unito.it

¹Regulation (EU) 2024/1689 of the European Parliament and of the Council of June 13, 2024, laying down harmonized rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance)

adaptability to accommodate future legal amendments and case law developments. Together, these components provide a pathway for understanding and navigating the intricate requirements of the Act, even as the Act itself continues to evolve. The DT-based framework is designed to address some of the basic issues in AI compliance by providing a systematic, repeatable classification methodology that guides AI providers through the legal requirements outlined in the AI Act. This framework can be situated within two main academic fields: in the field of legal visualization for regulatory compliance [4–6] and in the field of “Rule as Code” (“RaC”)² studies [8–10].

The main research question this paper addresses is “How can we build a reliable AI system automatic classification methodology for providers?”

This paper seeks to address the need for a simple yet solid methodology to tackle the issue of the AI Act classification through the use of a DT-based framework. This issue was initially analyzed in the context of the challenge “AI Act Compass: Navigating Requirements for High-Risk AI Systems” promoted by Legality attentive data scientists (LeaDS).³ The results of this experience have been presented in the workshop *CLAIRvoyant 2024* at the *JURIX 2024* Conference.⁴

The methodology of this paper builds on that experience and has been further developed in the project “EU Regulatory Sandboxes for AI” (EuSAIR). The Sandboxes are a “space for experimentation”⁵ prescribed by the AI Act where innovators can test the compliance of their systems with the AI Act. Stakeholders receive support in technical and legal compliance. The EuSAIR team is in charge of the legal analysis and audit for stakeholders. The DT is a choice made by our legal team to improve the readability of the AI Act, following the reading of legal professionals and translating it into a machine-readable format (DDL) and finally into an easy-to-use platform for non-domain-expert users.

This paper is structured as follows:

- Section 2 reviews the state of the art of rule-based decision-making systems and introduces Related Works;
- Section 3 introduces the chosen Methodology, examples of Norms’ Representation, Project Implementation, and a Case study;
- Section 4 presents some Limitations, Conclusions, and Future Work.

2. Rule-Based Decision-Making Systems

In this section, we review the current state of the art of rule-based decision-making systems, an established approach in AI and knowledge management for representing domain knowledge and supporting automated reasoning. A rule-based system is a type of expert system composed of a set of rules, typically

expressed in the form of if-then statements, which specify how given conditions lead to particular conclusions or actions. These systems range from relatively simple decision structures to more expressive formalisms capable of representing complex relationships and constraints [11]. This section explores prominent approaches to rule-based decision-making: DTs, Decision Model and Notation (DMN), Rule-Based Ontologies, DDL, and Answer Set Programming (ASP).

2.1. The naive approach

The if-then-else constructs represent the most basic form of rule-based logic, defining direct condition-action rules; they are the simplest way to implement a rule-based decision-making system. They are highly flexible and easy to implement, providing a foundation for rule-based reasoning that can be extended into more complex frameworks. However, as decision-making needs grow, nested if-then-else constructs quickly become unwieldy and hard to manage.

2.2. Decision modeling notation

The Decision Modeling Notation (DMN) is a standardized framework designed for managing business rules and decisions. It primarily uses Decision Requirement Diagrams and Decision Tables, enabling rule representation in a tabular format that’s highly interpretable and interoperable in business environments [12–14]. Decision tables capture complex if-then rules in a clear, consistent format that can be adapted to changes with minimal effort. However, DMN’s reliance on decision tables restricts its applicability for highly complex scenarios, such as those requiring dynamic rule deduction or handling interdependent, constraint-based logic. DMN lacks the depth needed for advanced reasoning in knowledge-intensive domains.

2.3. Decision tree

DTs represent hierarchical rule-based structures where decisions are made by recursively splitting data on specific attributes. Each node in a tree represents a condition, with branches representing the possible outcomes of that condition, leading to leaf nodes that define final decisions or classifications [15]. DTs are particularly popular in machine learning due to their high interpretability and ease of use in data-driven environments [16–18]. DTs are intuitive, making them ideal for rule-based tasks where interpretability is essential. But as data complexity grows, DTs can become deep, unwieldy, and prone to overfitting. Handling interdependent rules or constraints across nodes is challenging without advanced pruning or ensemble techniques.

2.4. Rule-based ontologies

Rule-Based Ontologies provide structured representations of domain knowledge, where entities, their attributes, and their relationships are defined according to ontological principles [19, 20]. Ontologies are commonly used in the Semantic Web, allowing rule-based inferencing based on well-defined relationships [21, 22]. Ontologies excel at organizing complex domain-specific information, but primarily model static representations [23–25]. Managing dynamic rules or handling complex logical interdependencies can be challenging, limiting their utility in high-stakes decision-making systems.

²Rules as Code is a label adopted by people in governments working on the idea of encoding legislation while it is being drafted, so that the logic of the resulting legislation can be read (and checked) by a computer, to improve the manner in which legislation is produced and the way in which it is available digitally [7].

³A European interdisciplinary project funded under the MSCA-ITN funding scheme available at <https://cordis.europa.eu/project/id/956562>

⁴Available at <https://jurix2024.law.muni.cz/>

⁵Article 57 of the AI Act

2.5. Rule-based DDL

DDL is an advanced form of logic that combines deontic logic (the logic of obligations, permissions, and prohibitions) with defeasibility. DDL allows rules to be overridden by exceptions or conflicting rules, providing a flexible framework for complex decision-making in environments where rules interact in nuanced ways. This makes DDL especially valuable for legal reasoning, policy compliance, and any domain where conflicting rules or exceptions are common. DDL excels in scenarios with interdependent and conflicting rules, allowing decisions to adapt dynamically based on contextual exceptions. Unlike other rule-based systems, DDL can process conflicting obligations (what must be done) and permissions (what may be done), allowing it to prioritize or dismiss rules as needed. This flexibility makes DDL an ideal choice for decision-making systems in highly regulated environments where conflicting rules, permissions, and exceptions are unavoidable. Implementing DDL requires specialized knowledge in logic programming and is computationally intensive, particularly when managing large rule sets. Its complex nature also makes it harder to interpret without additional layers for visualization or user guidance [26].

2.6. Answer set programming

ASP is a widely used logical language for knowledge representation and reasoning. Since its inception, the academic community has undertaken extensive research efforts, establishing ASP as one of the most widely adopted frameworks for AI, including applications at the industrial scale [27].

We decided to implement DDL using ASP as it allows us to define the rules, constraints, and facts of a designated problem. The problem is then solved by an ASP solver that finds a solution in the form of an answer set by grounding and finding the stable models [27].

2.7. Related works

The entry into force of the European AI Act has generated significant interest in methods and tools capable of supporting regulatory compliance.

Among the first operational compliance tools released by the European Commission is the *EU AI Act Compliance Checker*. The tool guides users through a sequence of binary questions to identify the obligations applicable under the AI Act according to their role (e.g., provider, deployer, importer, or distributor). Such interactive questionnaires provide an accessible entry point for organizations with limited legal expertise by helping them perform an initial assessment of the applicable regulatory requirements. Nevertheless, checklist-based approaches remain limited in terms of explainability and transparency. Although users may retrace the sequence of questions, the underlying legal reasoning and the specific provisions motivating the resulting classification often remain implicit. Consequently, these tools support compliance awareness but provide only limited assistance in understanding why a particular classification or obligation has been assigned.

Several academic works have proposed structured methodologies to simplify AI Act classification. Hanif *et al.* [28] present a question-based DT that guides users through successive questions to determine the risk category of an AI system. Their approach reduces the complexity of navigating the legislation and accelerates the classification process. However, users are intentionally shielded from the complete DT, limiting their understanding of the

overall regulatory logic. Furthermore, to the best of our knowledge, neither an interactive implementation nor the complete decision structure has been publicly released, limiting reproducibility and practical adoption.

Another major research direction relies on semantic technologies and knowledge representation. Golpayegani *et al.* [29] introduced AI Risk Ontology (AIRO), an ontology formalizing key concepts of the AI Act, including actors, risk categories, obligations, and lifecycle elements. Building upon this line of research, Hernandez *et al.* [30] propose an open knowledge graph that maps concepts and normative requirements between the AI Act and international standards. Their work addresses one of the central challenges of AI governance, namely, assessing the extent to which compliance with harmonized standards contributes to compliance with the regulation itself. By representing regulatory concepts and standard requirements as machine-readable knowledge graphs, the approach enables systematic comparison, reasoning, and traceability while facilitating future extensions as standards evolve.

Ontology-based approaches provide important advantages, including formal reasoning capabilities, automated querying through SPARQL, modular representations, and improved consistency in interpreting regulatory concepts. However, they also face several limitations. Legal obligations frequently depend on contextual interpretation, making them difficult to encode using static ontological structures. Moreover, maintaining such ontologies requires continuous updates as implementing acts, harmonized standards, and regulatory guidance evolve. Translating nuanced legal interpretations into formal semantic representations also remains a significant knowledge-engineering challenge.

Formal logic has likewise emerged as a promising approach for representing and reasoning over AI legislation. Lawniczak and Benz Müller [31] analyze the logical modalities embedded in the AI Act with the objective of enabling formal normative reasoning through the LogiKey framework. Their work investigates the suitability of higher-order logic embeddings for representing obligations, permissions, prohibitions, and other normative constructs within proof assistants such as Isabelle/HOL. Initial experiments demonstrate the potential of automated reasoning over regulatory provisions while also highlighting the complexity of capturing the semantics of legal norms in a computationally rigorous manner.

Related work on legal compliance checking extends beyond the AI Act itself. Francesconi and Governatori [32] propose reusable semantic patterns for modeling deontic norms and legal compliance within decidable fragments of OWL 2. Their framework represents obligations through ontology classes and property restrictions while supporting automated reasoning and defeasible norms. Although not specifically designed for AI regulation, these methodologies provide valuable foundations for machine-readable legal compliance systems and demonstrate how semantic technologies can support automated verification of normative requirements.

Recent research has also begun addressing specific compliance domains introduced by the AI Act. Gallese and Falletti [33] investigate compliance with the Regulation's data governance requirements for high-risk AI systems, emphasizing that fairness and bias mitigation should be addressed at the dataset level in addition to model behavior. Their work presents the FAnFAIR framework, which combines automated statistical analyses with qualitative assessments to evaluate dataset quality and fairness throughout the AI lifecycle. Rather than providing a general compliance framework, this approach focuses on one of the

most technically demanding aspects of the AI Act, namely, data governance and dataset assessment.

Overall, existing approaches primarily emphasize either interactive questionnaires, ontology-based knowledge representation, formal legal reasoning, or specialized compliance support for particular domains. Our framework shares the common objective of supporting AI system classification under the AI Act but adopts a complementary perspective. Rather than producing only a classification outcome, it explicitly models the decision process leading to each regulatory conclusion, making the underlying reasoning transparent and traceable. This improves explainability for both technical and non-legal stakeholders while providing a structured foundation that can support subsequent regulatory assessment, documentation, and compliance analysis.

3. Methodology

For our project, we decided to use DTs to provide visual support for nontechnical users in order for them to define rules. The rules have been selected from specific sections of the AI Act, and an ASP has been chosen as it provides flexibility for maintenance and to easily deal with exceptions to the defined rules, which can be changed over time as a legal framework.

3.1. AI Act selection

The AI Act version referred to in this paper is the latest version published on June 13, 2024. We were interested in isolating relevant articles of the AI Act related to the classification of an AI system and the resulting obligation for providers to have a clear compliance path for their roles.

Most of the AI Act focuses on high-risk systems, the criteria to include an AI in a high-risk class, the exceptions, the borderline cases, the requirements, etc. Thus, many of the articles that we decided to include are from Chapter III of the AI Act, the “HIGH-RISK AI SYSTEMS” chapter.

We decided to include 34 Articles, which vary from the scope of the AI Act to the permission to adopt voluntary codes of conduct, and 9 Annexes to give more details and specifications on obligations prescribed in the articles. Also, four recitals were taken into consideration as they give further information related to definitions of certain AI systems (e.g., general-purpose AI systems also known as “GPAI”) and exemptions from obligations (e.g., open-license GPAI that are not subject to transparency obligations).

Here is the AI Act selection:

- **Recitals:** 12, 54, 97, 104.
- **Articles:** 2, 3, 5, 6, 8, 9, 10, 11, 12, 14, 15, 16, 17, 18, 20, 21, 22, 43, 49, 50, 51, 53, 55, 60, 61, 72, 73, 95, 102, 103, 104, 105, 106, 107, 109.
- **Annexes:** III, IV, V, VI, VII, VIII, IX, XI, XII.

3.2. Decision tree

Once the Recitals, Articles, and Annexes have been isolated, our aim was to draft a model to visualize the pathway to compliance with the AI Act. We chose a DT model as it offers substantial advantages in compliance modeling due to its high transparency and interpretability.

The choice of which Recitals, Articles, and Annexes to be included is complex. The standpoint of the DT is the providers. Thus, there is no reference to distributors, importers, and other

roles in the AI development chain. Some articles in the AI Act refer to the Commission or the AI Office drafting new documents that providers might have to abide by (e.g., Art. 41 Common specifications). However, we gave priority to Articles prescribing obligations to providers, or in which providers are the main actors (e.g., Art. 43 Conformity).

The provisions considered in this analysis perform distinct normative functions: definitional norms (Recital 12; Arts. 2–3) establish the conceptual framework of the Regulation by defining its scope of application and clarifying the meaning of legal concepts, thereby ensuring interpretative consistency; classification norms (Recital 54, Articles 6–7, 51, 102–109, Annex III) determine the criteria according to which AI systems are allocated to different regulatory categories; prohibitive norms (Art. 5) categorically prohibit AI practices considered incompatible with fundamental rights or Union values, obligatory norms (Recitals 97 and 104, Arts. 8–22, 43, 49, 50, 53, 55, 60, 72, and 73) impose substantive duties on providers belonging to different classes of risk; specification norms (Annexes IV–XII) complement these obligations by detailing the technical, procedural, and documentary requirements through which compliance must be demonstrated, including conformity assessment procedures and technical documentation. Finally, Article 95 has been included as a recommended measure for AI systems that do not fit in any other category.

Each branch in the tree illustrates a clear, logical decision pathway, which allows stakeholders, including those without technical expertise, to properly understand and manage the Regulation. DTs’ rule-based structure aligns with the requirements of compliance models, and this inherent structure facilitates the capture, audit, and validation of compliance requirements, making it a natural fit for regulatory modeling. The EuSAIR team is responsible for the legal analysis and audit activities supporting stakeholders throughout the AI system classification process. The DT represents a methodological choice developed by the legal team to improve the readability and practical understanding of the AI Act, translating its legal structure into a more accessible and operational framework.

The sequence of questions within the DT framework follows the logical flow that legal professionals would typically adopt during an audit process with stakeholders. In particular, the DT reflects the structure of the AI Act itself: the assessment begins by determining whether the system falls within the scope of an AI system as defined by the regulation, followed by an evaluation of its characteristics, complexity, and operational context. The analysis then proceeds through questions concerning the type of model employed, the intended purpose, the deployment context, and the nature of the interaction between the AI system and users, which are essential elements for determining the applicable risk classification. The DT also incorporates the specific exceptions and exclusions associated with each category, ensuring that the classification process accurately reflects the regulatory conditions established by the AI Act. This approach allows stakeholders to follow a structured assessment pathway that mirrors both legal reasoning and the regulatory logic of the AI Act.

Our DT has been modeled in Business Process Model and Notation as it is a standardized graphical notation that depicts the steps of a chosen process [34–37]. The business process modeling and management software platform used to visualize processes is iGraphix.

In Figure 1, the process is shown from a high-level perspective. On that premise, the DT presents four steps: scope, model,

context, and interaction. Based on this division, there are four general questions to be answered.

The questions to answer are:

- Is it in the scope? The AI Act gives some grounds of exception for specific AI systems or uses.
- Which AI model? The AI Act gives information and subsequent obligations based on models and their computational power.
- In which context? In the AI Act, contexts define most of the obligations to follow.
- Is it interacting with people? Lastly, some obligations are related to the human–AI relationship.

All the questions are organized in a structured progression to serve a dual purpose: first, it reduces redundancy by ensuring that each stage of the assessment builds upon prior determinations, and second, it creates an increasingly focused analytical pathway, moving from broad assessments of applicability and scope toward the precise evaluation of context-specific regulatory criteria.

In Figure 1, each question has a set of sub-processes that give further details on the compliance path; the cross inside each activity box symbolizes the presence of a sub-process. Each will be discussed below.

As for the first question, the DT begins by distinguishing between systems that fall within and outside the regulatory purview of the AI Act, as shown in Figure 2. This initial

Figure 1
Representation of the Decision tree main process





Figure 2
Representation of the first sub-process of Figure 1: Scope




classification ensures that any AI system falling under exemptions (e.g., military or third-country authorities) is recognized and excluded from compliance requirements at the earliest possible stage. This approach minimizes unnecessary processing for systems exempt by design, aligning with the principle of regulatory efficiency. The white circle in the sub-process means “begin event” in case it is at the start of a process or a sub-process; it means “event ended” in case it is at the end of a process or a sub-process; the gray circle inside the circle means “process ended.”

In Figure 3, the sub-process defines which models are AI systems and whether the system can be classified as a general-purpose AI (GPAI). GPAIs are defined by their adaptability across diverse applications and large-scale data processing capabilities. Within this framework, the tree differentiates between open GPAI (those open source, free license, and giving substantial information regarding its structure) and general GPAI classifications, where open GPAIs do not require transparency obligations due to their open-access and broad availability characteristics. The inclusion of model-related questions ensures that providers comprehensively address obligations related to model adaptability and information transparency. Finally, the AI Act has established additional obligations for “GPAIs with substantial risks” and has set the computational threshold of 10^{25} floating-point operations (FLOPs).

Figure 4 is the core of the AI Act DT; it analyzes the context of the use of an AI system to assess the number of obligations that the provider will have to abide by. It is built on features that are linked to the Banned Class of the AI Act and on the High-Risk Class.

- It starts with “checking for banned” as, in case an AI system is classified as banned by the AI Act, it cannot enter the EU market, be developed inside EU territories, or even have its output impacting data that goes through the EU data space. Providers should be aware that their tools are completely out of the market. However, as can be seen in the first part of Figure 4, even banned practices have exceptions that allow their use. Each of the banned categories in the AI Act has specific features that possibly “lower” their classification.
- Thus, the analysis carries on to “checking for high-risk.” Here the process splits into three paths that can lead to the classification of high-risk. These three paths correspond to three different norms inside the AI Act: AI tools already under New Legislative Framework (NLF) legislation (Article 6 and Annex I section A); AI tools with specific uses (Annex III); and AI tools under other Union harmonization legislation (Annex I section B).

Figure 3
Representation of the second sub-process of Figure 1: AI model



AI tools already under NLF legislation also have to undergo a third-party assessment under their own framework to be considered high-risk, whereas AI tools under Annex III have to be used for main tasks/not irrelevant tasks to be considered high-risk. For example, preparatory tasks are irrelevant under the AI Act. AI tools under other Union harmonization legislation follow another specific set of norms (Articles 102–109) because of their peculiar nature. It is relevant to note that testing in the real world adds other layers of Articles, as it is also considered a risk to take. Finally, it is highlighted that if the provider is a small-medium enterprise (SME), its compliance procedure is simplified.

The last sub-process in the DT is in Figure 5. It focuses on systems involving direct user interaction or the generation of synthetic content (e.g., images, audio). This sub-process addresses limited-risk classifications, where the primary obligations pertain to transparency and informed user consent. For example, if an AI system generates synthetic media or interacts directly with users (such as a chatbot), providers are required to disclose the synthetic nature of content or the presence of an AI system in user interactions.

For systems that do not fall into any of the defined classes (banned, high-risk, or limited), the AI Act prescribes a voluntary code of conduct as the residual compliance measure. This

Figure 4
Representation of the third sub-process of Figure 1: Context





Figure 5
Representation of the fourth sub-process of Figure 1: AI interaction with persons



voluntary approach allows providers to demonstrate proactive adherence to best practices, even in cases where formal regulatory obligations are minimal.

3.3. Defeasible Deontic Logic

In this subsection, we introduce the foundational concepts of DDL.

DDL is a logical framework that integrates **deontic logic** with **defeasible reasoning**, which allows conclusions to be retracted in the face of new evidence or stronger rules. We used the ASP implementation of DDL of [38].⁶

Deontic Logic deals with normative statements about what is obligatory, permitted, or forbidden. Traditional deontic logic treats norms as absolute, lacking mechanisms to handle exceptions or conflicts.

Defeasible Reasoning introduces rules that can be overridden by other rules, enabling reasoning that mirrors real-world scenarios where exceptions and priority among rules are common.

Atoms are the basic, indivisible propositions or statements that serve as the fundamental units of meaning in the logical system. They are used to build more complex expressions through logical connectives and deontic operators, such as obligations (O), permissions (P), and prohibitions (F).

By combining these, DDL allows us to model complex normative systems where rules can have exceptions, and conflicts between norms can be resolved based on contextual priorities.

Implementation in Answer Set Programming (ASP)

ASP is a declarative programming paradigm suited for non-monotone reasoning, making it ideal for implementing DDL. ASP allows us to encode rules, facts, and priority relations declaratively and to compute the logical consequences efficiently.

Our implementation comprises the following key components (for an in-depth explanation of the implementation, see [38]):

Representation of facts and rules

- **Facts:** basic assertions about the current state of the world
 - **Constitutive Rules:** fundamental norms that are generally applicable unless overridden: `constitutiveRule(R, X). applicable(R, X):- Conditions.`
 - **Prescriptive Rules** establish obligations: `prescriptiveRule(R, X). applicable(R, X):- Conditions.`
 - **Permissive Rules** establish permissions: `permissiveRule(R, X). applicable(R, X):- Conditions.`
- Where **applicable** defines under which conditions rules apply often involving defeasible predicates. `applicable(R, X):- defeasible(Condition1), defeasible(Condition2).`

Rules are subject to a hierarchy, defined by the **superior** predicate in order to resolve conflicts. The derivation for obligations and permissions is straightforward from **prescriptiveRule** and **permissiveRule** predicates.

Defeasibility and Conflict Resolution

- **Defeasible Conclusions.** A conclusion X holds defeasibly if there is no opposing fact and there is an applicable constitutive rule for X that is not overruled: $\text{defeasible}(X)$:- $\text{opposes}(X, X1)$, not $\text{fact}(X1)$, $\text{constitutiveRule}(R, X)$, $\text{applicable}(R, X)$, not $\text{overruled}(R, X)$.
- **Overruling Rules.** A rule R is overruled if there is an applicable and undefeated constitutive rule $R1$ for an opposing literal $X1$: $\text{overruled}(R, X)$:- $\text{opposes}(X, X1)$, $\text{constitutiveRule}(R, X)$, $\text{constitutiveRule}(R1, X1)$, $\text{applicable}(R1, X1)$, not $\text{defeated}(R1, X1)$.
- **Defeating Rules.** A rule R is defeated if there is a superior rule $R2$ for an opposing literal $X2$ that is applicable: $\text{defeated}(R, X)$:- $\text{opposes}(X, X2)$, $\text{constitutiveRule}(R2, X2)$, $\text{superior}(R2, R)$, $\text{applicable}(R2, X2)$.

⁶The implementation is available at <https://github.com/gvdgdo/Defeasible-Deontic-Logic>.

Where the **opposes** predicate represents a conflict between two literals.

Deontic conflict resolution

There are two ways to resolve conflicts in DDL: one is by overruling and the other by defeating:

- **Deontic Overruling.** An obligation or permission is overruled if there is an applicable, undefeated deontic rule for an opposing action: `deonticOverruled(R, X):- prescriptiveRule(R, X), deonticRule(R1, X1), opposes(X, X1), applicable(R1, X1), not deonticDefeated(R1, X1).`
- **Deontic Defeating.** A deontic rule is defeated if there is a superior opposing deontic rule: `deonticDefeated(R, X):- opposes(X, X1), deonticRule(R1, X1), applicable(R1, X1), superior(R1, R).`

We choose to operate on both DTs for better visualization for nontechnical users and DDL for better expressiveness of the rules. The conversion between a DT module and the corresponding DDL was made by hand. Here are some examples:

Banned uses

Here we have translated from the DT the part regarding the forbidden uses of AI by the AI Act into a rule.

In Figure 6a, we can see how the four different banned uses in 6a become four different atoms in 6b that makes the **banned_use** atom positive. Moreover, the exceptions are also translated, which is one of the advantages of using DDL, where if certain conditions that allow the use of the AI are met, the **banned_use** atom becomes negative.

First of the high-risk routes

In the AI Act, we have defined three main routes to define high-risk systems and their related obligations. These routes, as described before, are linked to Article 6, Annex I Section A, Annex I Section B, and Annex III. The route that we show here is related to Article 6 and to AI systems already under other EU regulations (“NLF” regulation).

In Figure 8, we can see how the 12 different banned uses in 8a become 12 different atoms in 8b that makes the **high_risk_article6_1** atom positive, which represents only the part

of scenarios represented in 6a, only if the **third_party_assessment** is also positive.

Each atom shown in both Figures 6b and 8b is a basic proposition of rules, and each of them is linked to an implicit question. As an example, the atom “**criminal_offence_prediction**” is linked to “Is it used for Criminal Offense prediction?” In the same way, the high-risk route is shown in Figure 8b; the atom “**machinery**” embeds the implicit related question “Is it included in the Machinery Directive?”

The most relevant questions have been expressed in the DT, whereas others have been deemed less critical and therefore omitted to maintain clarity in the workflow.

Also, Figures 7 and 9 are the prescriptive rules that embed obligations for each and every atom. In these examples, obligations related to banned uses of AI are mainly expressed in Article 5. For the high-risk route, it is evident that the number of obligations has multiplied in the related Articles. This mainly happens because high-risk AI systems are the one that bears the most requirements to comply with to enter the EU market, whereas banned use AI cannot enter the EU market whatsoever. This means that they bear just one obligation: to be banned.

3.4. Transition between the DT framework and DDL implementation

The DT framework has been chosen so as to ease the visualization of the AI Act: the division into four sub-processes served the purpose of dividing the main topics of the AI Act, namely, being in the scope of the AI Act, being a GPAI, the context of use, and belonging to a residual class (limited or minimal) based on their interaction with people. The steps of this framework have been organized into questions for the user to guide them through AI system compliance. The content of these questions, and consequently their correctness, is derived from the conditions of the DDL rules, specifically from the atoms appearing in the antecedent of each rule and from the correspondence between the structure of the DDL rules and the structure of the relevant legal provisions.

⁷Directive 2006/42/EC

Figure 6
Representation of banned uses in the decision tree and in the DDL. (a) The banned uses in the decision tree; (b) the banned uses in a DDL rule




Figure 7
Prescriptive rule of banned uses

```
prescriptiveRule(t25o,article5).
applicable(t25o,article5) :-
    defeasible(ai_scope),
    defeasible(ai_system),
    defeasible(untarget_scraping),|
    defeasible(banned_use).

hasSource((t25;t25o),
    (article5;article5_1_d;article5_1_e;
    article5_1_f;article5_1_g;article5_1_h)).
```

Figure 8
Representation of the high-risk route in the decision tree and in the DDL. (a) The first of the high-risk routes in the decision tree; (b) the first of the high-risk routes in a DDL rule




Figure 9
Prescriptive rules of the first of the high-risk route

```
prescriptiveRule(t38_o,
    (article8;article9;article10;article11;article12;article13;article14;
    article15;article16;article17;article18;article19;article20;article21;
    article22;article49_1;article72_4;annex_iv;annex_v;annex_vii)).
applicable(t38_o,
    (article8;article9;article10;article11;article12;article13;article14;
    article15;article16;article17;article18;article19;article20;article21;
    article22;article49_1;article72_4;annex_iv;annex_v;annex_vii)) :-
    defeasible(high_risk_article6_1).
```

The transition between the DT framework and DDL serves another purpose: formalizing the framework in a machine-readable format. The transition between the two is manual: the legal team has collaborated with IT professionals and logic domain experts to perform such conversion.

Converting legislation into machine-executable code requires more than programming; it combines legal interpretation with technical implementation. As part of this process, we adopted the method of technical coding validation [39], whereby an encoded provision is considered valid when it adheres to the selected

coding language and relevant coding conventions. Technical validation ensures that the code follows programming language rules, uses consistent naming conventions, and integrates correctly with other encoded rules. To address these challenges, we have relied on legal alignment to ensure the encoded rules accurately reflect the legislation's text and purpose. This process also acknowledges that some legal concepts remain inherently ambiguous or require human judgment and therefore cannot be fully resolved through programming. Consequently, successful encoding of legislation depends on interdisciplinary collaboration between

legal professionals and programmers to produce accurate, reliable, and maintainable machine-readable representations of legislation. This was made possible by developments in the context of the EuSAIR team project.

Furthermore, the DDL implementation is structured around the tasks of the *DT framework*, where each task represents a discrete step in the interpretation and formalization of regulatory provisions.

For example, in Figure 6b, Task 25 (t25) is represented. It goes from the beginning of the third sub-process in Figure 4 to the closing gateway in Figure 6b. It formalizes the prohibition of specific AI practices under Article 5 of the AI Act.

A *constitutiveRule* defines a legal concept or status, in this case, *banned_use*, while the applicable predicate specifies the conditions under which that concept holds. These conditions are expressed through defeasible predicates, which represent legal facts or assumptions that may be rebutted or overridden if contrary evidence or legal exceptions arise. In the example, *banned_use* is established if at least one prohibited practice (e.g., *realtime_biometric* or *criminal_offence_prediction*) is defeasibly satisfied. Similarly, the prescriptive rule *t25o* becomes applicable only when all required conditions are met, including the existence of an AI system within the scope of the regulation and the previously established *banned_use*. Finally, the *hasSource* predicate ensures traceability by linking each task and its associated rules to the corresponding legal provisions. Since a single task may derive from multiple legal sources, *hasSource* can reference several provisions simultaneously, including articles, paragraphs, points, and subpoints (e.g., Article 5, Article 5(1), points (d)–(h)), thereby preserving the correspondence between the formal representation and the original legislative text.

The following section presents the implementation of the proposed platform, which enables stakeholders to determine their system classification and the corresponding regulatory obligations. It describes the practical realization of the project and the functionality that will be available to end users.

3.5. Project implementation

Our Project outcome is a user-friendly platform called “Compl-AI.” For Compl-AI, the technical implementation was constructed around a core set of Python tools, Flask for the web interface, Clingo for processing ASP rules, and various supportive libraries to enable interaction, data handling, and logic processing. Here’s a more detailed breakdown:

Technical Implementation

To implement DDL, we used Clingo, a tool for ASP, used to interpret and solve the rule-based logic for the application.⁸ Python’s *clingo* library provided an interface for executing ASP programs within a Python environment, allowing the system to dynamically load DDL rules and query defeasible logic scenarios.

As for the web interface, it was built with Flask, which allowed the system to serve a user-friendly, browser-based interface. Flask’s routing and request-handling features facilitated user input and query processing.

Application Workflow

The users will input rules or queries through the web interface, which Flask routes to the backend. The backend logic takes

these inputs, formats them according to DDL syntax, and then uses Clingo to evaluate the rule set.

Clingo then determines the appropriate response based on the ASP output and structures this response for the user. Finally, Flask renders the results on the web interface, providing users with a clear summary of rule evaluations. Additionally, if the system detects conflicts or overrides (a feature of defeasible logic), it flags these to ensure users understand any nuanced reasoning outcomes.

This setup enabled Compl-AI to process complex rule-based queries in real-time and offer a web interface accessible to non-technical users while leveraging the power of Clingo’s ASP engine for defeasible logic processing.

3.5.1. Compl-AI Web Interface

The functionality starts with the user interacting through a straightforward web interface. Upon accessing the system, users provide details about their AI system by responding to a series of boolean questions. These questions are divided into five distinct sections: Scope, Banned Risk, Context, Interaction with People, and Phases. By organizing the questionnaire in this manner, users can systematically specify critical attributes of their AI system, such as its intended domain, risk category, deployment environment, expected user interaction, and lifecycle stage. The system then aggregates these responses to form a set of facts about the AI system.

Behind the scenes, the backend takes these facts and applies them against our predefined logic rules to determine which regulatory requirements are triggered. Our rules database is composed of a dynamic set of Direct Description Logic (DDL) statements that encapsulate various AI-related regulations. Once the relevant facts are extracted from the user’s input, the backend processes them through an ASP engine, specifically Clingo, to produce an output of the regulations applicable to that particular AI system. Clingo’s nonmonotonic reasoning capabilities are instrumental here, as regulations often overlap or can come into conflict, and the solver’s approach is particularly well-suited for navigating these complexities.

For the user interface, we chose the Flask framework. Flask’s lightweight nature and modular design allowed us to create user-friendly forms that capture the details of the AI system and then display the final results in a clear, organized manner. Flask’s extensive ecosystem also made it possible to integrate seamlessly with our backend logic and to provide a responsive, intuitive dashboard. Users can view which laws, directives, or guidelines are considered relevant for their AI system, along with any critical notes or considerations about compliance.

In terms of architecture, we deliberately kept the tool’s design modular to ensure maintainability and extensibility. The web interface layer, developed in Flask, functions independently from the core logic layer, which houses the DDL rules and handles interaction with the Clingo solver. This separation means that modifications or additions to regulatory rules do not require changes to the interface, and vice versa. Moreover, any future modification or expansion in the rule base, such as adding support for new AI regulations, is greatly simplified by this decoupled design. We can continuously update our logic to reflect emerging legislative trends and integrate new reasoning capabilities as necessary.

In addition to modularity, our approach emphasizes transparency and traceability. For instance, if a user wants to understand why a particular regulation appears in the dashboard, we provide a brief rationale based on the logic rules used. This level of detail

⁸<https://github.com/potassco/clingo/releases>

not only helps clarify the results for nontechnical audiences but also supports auditors, legal professionals, and other stakeholders seeking an explanation of compliance requirements.

Finally, we designed our tool with practical usage in mind. By providing a guided questionnaire and automatically generating a set of applicable regulations, we aim to reduce the complexity and effort required for compliance activities. This is especially important in rapidly evolving fields such as AI, where new guidelines and governance frameworks emerge frequently. Our architecture and workflow are intended to remain adaptable, ensuring that organizations can rely on the tool as a cornerstone of their compliance strategy both now and in the future.

A Case Study from the EuSAiR Project

This section will present one case study from the EuSAiR Project. The EuSAiR Project (EU Regulatory Sandboxes for AI) is a two-year initiative funded by the EU's Digital Europe

Programme to support the implementation of AI Regulatory Sandboxes across all 27 EU Member States, as required by the EU AI Act. These sandboxes provide controlled environments where AI developers can safely test and validate innovative AI systems under regulatory supervision before market deployment. The project develops common guidelines, tools, and benchmarks to harmonize sandbox operations, prioritizes access for SMEs and startups to reduce compliance costs, and equips national authorities with operational standards, risk assessment frameworks, and training resources. EUSAiR also gathers best practices and real-world evidence to inform future AI policy and conducts pilot testing on European high-performance computing infrastructures, while collaborating with European Digital Innovation Hubs and Testing and Experimentation Facilities to build a coordinated AI innovation ecosystem.

In this context, one model will be examined: an AI-powered algorithm that monitors patients' sleep patterns. The users we refer to are the developers of the AI system.

Figure 10
Tool's main page




Figure 11
Scope selection questionnaire page

Compl-AI: Scope

- ☐ Is the AI used for military purposes?
- ☐ Is the AI used by third country authorities?
- ☐ Is the AI used for Research and development purposes

Does it employ one or more than one of these techniques?

- | | |
|---|--|
| ☐ Linear regression methods and logistic regression methods
(e.g. machine learning to approximate function, to optimize problems, to maintain or improve performance) | ☐ Simple prediction system
(e.g. basic statistical learning rule, static estimation system) |
| ☐ Machine learning models for prediction and basic data processing
(e.g. describing info, for testing hypothesis, for data visualization) | ☐ Natural language processing |
| ☐ Classical heuristics
(e.g. trial-and-error strategy, pattern recognition, data-driven learning) | ☐ Other machine learning approaches
(e.g. supervised learning, reinforcement learning, unsupervised learning, self supervised learning, deep learning) |

Upon opening the web interface, the users are presented with the main page, and they click on “Start Form” to start the assessment (see Figure 10). For each question, the cube on the left should be ticked to indicate “yes” and left unticked to indicate “no.”

In the Scope section (Figure 11), they are faced with multiple choices regarding the possible exclusion of the system from the AI Act or whether the system falls under the GPAI. This is the web interface equivalent of Figures 2 and 3. A few more options related to the AI model techniques have been implemented under the “Does it employ one or more than one of these techniques?” This choice was due to the publication of the guidelines on the definition of AI systems⁹ occurring after the first draft of the DT framework. These updates have been directly added to the rules and did not cause any change in the DT framework (see Figure 12).

Since this system does not fall under any of these assumptions related to exclusions from the AI Act and GPAI, the users did not select anything and clicked on Next.

In the Banned Risk section (Figure 13), there are multiple choices related to practices banned under the AI Act. The web interface also takes into consideration the related exception to each category of banned practices.

The system of reference does not fall under any category, so the users did not select anything and clicked on next.

In the Context section (Figure 14), they will be faced with multiple choices regarding the main aspect of the system itself, for instance, the context in which it will be applied. Since the system we are taking into consideration is a medical one, the users have chosen “Regulation on Medical Devices”; we also selected “Intended by the manufacturer to have a medical purpose” and ticked “Does it already require third-party assessment?”

Indeed, the AI system is a system already introduced in the market and sold as a medical device to monitor patients’ sleep patterns. Finally, the users clicked on next.

In the Interaction with People section (Figure 15), they are faced with multiple choices regarding how the people will interact with the system if the system interacts with people. Here, the users ticked “Is it actively interacting with people (e.g., chatbot)?” since the system is taking data actively from people and clicked on next.

In the Phases section (Figure 16), the stage of development of the system can be chosen. This section has been created to select the right amount of obligations needed at each stage. In general, a provider will be interested in the total amount. This option, however, allows users to check obligations step by step. This section has been created to connect the current work with future research. It will be developed to be linked to our Handbook “AI Compliance Framework: A Handbook For Compliance Tools based on AI Acts’ Requirements” [40]. In the handbook, each obligation is divided into three phases: Compliance (Project phase), Conformance (Development & Deployment), and Audit (Monitoring). This division is designed to underline in detail practical actions to be taken in various phases of the creation of an AI system.

In this specific example, the users have ticked “Is your AI system in surveillance/post-market phase?” and clicked on Next.

This page presents the results of the system classification process (Figure 17). In the considered case, the system has been classified as both a high-risk and a limited-risk AI system. The highest applicable classification represents the primary reference category, while the coexistence of multiple classifications indicates that the obligations associated with each relevant risk category must be considered and addressed.

The applicable obligations can be consulted through three complementary sections:

- **“Your Obligations Summarized” section:** This section is intended for stakeholders seeking an overview of the main requirements arising from the AI Act. The relevant provisions have been summarized by the legal team of EuSAIR, highlighting the fundamental principles and regulatory objectives underlying the corresponding Articles without reproducing their full text.
- **“Full Articles” section:** This section provides access to the complete text of the relevant AI Act provisions, enabling users and stakeholders to consult the original regulatory sources in their entirety.
- **“Other Applicable Regulations” section:** This section identifies additional regulatory frameworks that may apply to the system, providing stakeholders with a broader overview of the compliance requirements that must be considered alongside the AI Act.

In the considered case, the system described as an “AI-powered algorithm that monitors patients’ sleep patterns” has been classified as both a high-risk and a limited-risk AI system. This

⁹Commission Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689 (AI Act)

Figure 12
Updated details, referring to the Scope sub-process in Figure 2

```
constitutiveRule(t2_2,non(ai_system)).
applicable(t2_2,non(ai_system)) :-
    defeasible(
        rule based;
        statistical_model;
        automation_and_scripting;
        control_system;
        static_algorithm;
        embedded_system;
        simulation_and_modelling;
        basic_recommendation_system;
        data_processing_tool;
        traditional_software_application;
        others_coming_soon)
```




Figure 13
Banned Risk questionnaire page

Compl-AI: Banned Risk

- ☐ Does it work on criminal offence prediction?
 - ☒ - Is it for support and based on previous offences?
- ☐ Is it used for Emotion inference in education\workplace?
 - ☒ - Is it used for Medical or safety reasons?
- ☐ Does it use biometric data to infer sensitive attributes (race, political opinions, etc.)?
 - ☒ - Is it lawfully acquired?
- ☐ Is it a real-time biometric identification system for law purposes?
 - ☒ - Is it intended to be used by authorities for Locating Crime Victims, Preventing Imminent Threats or Identifying Serious Crime Suspects
 - ☒ - Is there an authorization by a judicial authority?

Figure 14
Context questionnaire page

Compl-AI: Context

- ☐ Is it a biometric systems?
 - ☒ - Is it used only to confirm a self claimed identity?
 - ☒ -- Has it been\will it be tested in real-word conditions?
 - ☒ - Does it carry out a narrow\irrelevant task (e.g. preparatory task)?
- ☐ Is it used for the management and operation of critical digital infrastructure?
 - ☒ - Has been\will be tested in real-word conditions?
 - ☒ - Does it carry out a narrow\irrelevant task (e.g. preparatory task)?
- ☐ Is it used to manage education and vocational training, such as access to education\training ,evaluation in education\training, assessment of individuals, monitoring and detecting behaviour?
 - ☒ - Does it carry out a narrow\irrelevant task (e.g. preparatory task)?

classification implies that the system must comply with the full set of obligations associated with both risk categories. Therefore, the applicable requirements include Articles 8, 9, 10, 11, 12, 14, 15, 16, 17, 18, 20, 21, 22, 43, 49, 50, 60, 61, 72, and 73 of the AI Act, as well as the relevant requirements established in Annex IV, Annex V, and Annex VII.

Among these obligations, Article 50 represents the specific requirement associated with the limited-risk classification. Consequently, while the high-risk classification defines the main compliance requirements, the limited-risk classification adds additional transparency obligations that must also be satisfied.

4. Limitations, Conclusion, and Future Works

The proposed multi-choice compliance platform provides AI providers with direct access to structured information, guidance, and clarifications regarding AI classification and regulatory

obligations under the AI Act. By adopting a rule-based modeling approach rather than relying on large language models or other generative AI technologies, the platform offers several advantages, including full explainability, deterministic reasoning, lightweight implementation, and complete traceability of every step in the compliance path. Furthermore, the COMPL-AI framework structures the provisions of the AI Act into actionable decision pathways, enabling regulatory requirements and supporting documents to be systematically linked and adapted as additional European and national guidance becomes available.

Despite these advantages, the rule-based approach presents inherent limitations. As widely recognized in the literature, rule-based systems face challenges related to scalability, computational complexity, maintainability, adaptability, and knowledge acquisition [41]. As the number of regulatory provisions, exceptions, and sector-specific requirements increases, the corresponding rule base inevitably grows in size and complexity, making

Figure 15
Interaction with People questionnaire page

Compl-AI: Interaction with people

- ☐ Is it actively interacting with people(e.g.chatbot)?
- ☐ Does it generate synthetic content?
- ☐ Is it just used for editing?

NEXT

Figure 16
Phase selection questionnaire page

Compl-AI: Phases

- ☐ Is you AI system in project phase?
- ☐ Is you AI system in development phase?
- ☐ Is you AI system in surveillance/post market phase?

NEXT

consistency management, redundancy detection, and rule maintenance increasingly demanding [42]. Moreover, the current implementation relies on the manual formalization of legal provisions into decision rules. Although this approach is inherently resource-intensive and requires continuous maintenance to reflect legislative developments, it has been intentionally adopted to preserve legal accuracy, semantic consistency, and full traceability of the resulting compliance pathways. The primary advantage of employing DDL lies in its ability to represent open-textured legal terms, exceptions, and potential conflicts between rules in an explicit, structured, and monitorable manner. Furthermore, several empirical validation instances have demonstrated that the resulting encoded rules generally correspond to the interpretations provided by human users, as they provide conceptual structures that closely align with the requirements of legal texts [8].

Also, the formalization and validation of the rule base were carried out through a multidisciplinary collaboration involving legal experts, IT professionals, and specialists in deontic logic, providing complementary expertise to ensure that both the legal interpretation and the formal representation of regulatory requirements are robust and consistent. It is also important to note that the objective of this work was not to develop a static representation of the current AI Act but rather to establish a solid and extensible compliance framework. At the time of writing, the AI Act has been formally adopted, while many implementing acts, harmonized standards, and official guidelines are still under development. Consequently, instead of attempting to model provisions that are not yet available, this work focuses on defining a structured compliance methodology, identifying appropriate compliance analysis categories, and designing

a DT capable of accommodating future regulatory developments without requiring substantial redesign.

From a reasoning perspective, rule-based systems are primarily designed to capture explicit and deterministic knowledge and therefore exhibit limited flexibility when dealing with uncertain, context-dependent, or evolving regulatory scenarios. This limitation is particularly relevant for normative reasoning. Governatori [43] demonstrates that rule-based systems adopting well-founded, grounded, and skeptical stable semantics are unable to capture *weak permission* in the presence of unresolved deontic conflicts. When two conflicting obligations coexist without a mechanism for resolving precedence, these semantics refrain from deriving either obligation and consequently fail to infer the corresponding weak permission, even though weak permission is conceptually defined as the absence of an obligation to the contrary. Consequently, purely rule-based compliance systems may not fully represent complex legal reasoning where conflicting obligations, exceptions, or interpretative ambiguities exist.

Nevertheless, the modular architecture of the proposed platform mitigates several of these challenges by allowing incremental updates and extensions without redesigning the overall framework. Rather than representing a fixed sequence of regulatory questions, the DT has been designed as a modular structure whose branches can be independently extended to incorporate additional regulatory domains. For example, what currently appears as a single decision task regarding whether an AI system qualifies as a medical device can naturally evolve into a dedicated parallel decision pathway that guides users through the Medical Device Regulation compliance process. The same architectural principle can be applied to other sector-specific legislation within the NLF, enabling the platform to support integrated multi-regulation compliance while preserving consistency and traceability across different regulatory domains.

The dynamic nature of the European regulatory landscape further reinforces the need for continuous evolution of the platform. Additional implementing acts and guidance from the European Commission are expected, including documents supporting high-risk AI requirements, harmonization measures, simplified conformity assessment procedures, and potential alignment with international standards such as ISO. Rather than viewing these updates as limitations, the proposed framework has been designed to accommodate them through controlled revisions of the rule base and decision pathways, ensuring that regulatory changes can be incorporated while preserving the consistency and auditability of the compliance model.

Future work will therefore focus on continuously monitoring and incorporating regulatory updates as new implementing acts and official guidance become available. The modular structure will be extended with sector-specific decision pathways, richer traceability between legal provisions and formalized rules, and mechanisms supporting the impact analysis of regulatory amendments. Additional user-oriented functionalities, including customized compliance workflows, role-specific compliance checklists, and semi-automated support for maintaining the rule base, will also be investigated. Although the current maintenance process is entirely manual, this choice ensures rigorous expert validation during this initial phase. As the regulatory framework matures, semi-automated knowledge acquisition and hybrid reasoning techniques will be explored to improve scalability while preserving the explainability and traceability that constitute the primary strengths of the proposed approach. Future work will also include further validation of the proposed framework. The present paper primarily aims to introduce the methodology and

Figure 17
Applicable regulations page

demonstrate its functionality through its application to the legal assessment of the AI Act. Further empirical validation, based on additional case studies, will be presented in subsequent research. These studies will further investigate the framework's accuracy, reliability, scalability, and overall effectiveness.

In conclusion, this work presents a structured, explainable, and extensible compliance platform that assists AI providers in navigating the regulatory obligations introduced by the AI Act. Rather than proposing a compliance solution limited to the current legislative state, the main contribution lies in the definition of a robust compliance structure based on well-founded analysis categories, explicit decision pathways, and complete traceability between legal provisions and formalized rules. This architecture provides a stable foundation upon which future European guidance, implementing acts, harmonized standards, and sector-specific regulations can be progressively integrated. Although rule-based modeling introduces recognized limitations, particularly regarding maintenance effort and the representation of complex normative conflicts, its transparency, auditability, deterministic behavior, and expert-validated knowledge representation make it particularly well-suited for regulatory compliance applications where explainability and legal accountability are essential.

Declaration on Generative AI

Generative artificial intelligence tools were used for language editing and spell-checking. The authors used OpenAI ChatGPT-4 and Grammarly for such purposes. They remain fully responsible for the content and intellectual contributions.

Ethical Statement

This study does not contain any studies with human or animal subjects performed by any of the authors.

Conflicts of Interest

Iliara Angela Amantea is the Peer Review Board Member for *Journal of Computational Law and Legal Technology* and was

not involved in the editorial review or the decision to publish this article. The authors declare that they have no conflicts of interest to this work.

Data Availability Statement

Data available on request from the corresponding author upon reasonable request.

Author Contribution Statement

Marinella Quaranta: Conceptualization, Methodology, Software, Validation, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization, Supervision, Project administration. **Alessandro Bucci:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Visualization, Supervision, Project administration. **Iliara Angela Amantea:** Conceptualization, Validation, Resources, Writing – original draft, Visualization.

References

- [1] Quaranta, M., Amantea, I. A., & Grosso, M. (2023). Obligation for AI systems in healthcare: Prepare for trouble and make it double? *The Review of Socionetwork Strategies*, 17(2), 275–295. <https://doi.org/10.1007/s12626-023-00145-z>
- [2] Guldemann, P., Spiridonov, A., Staab, R., Jovanović, N., Vero, M., Vechev, V., & Vechev, M. (2024). COMPL-AI framework: A technical interpretation and LLM benchmarking suite for the EU Artificial Intelligence Act. *arXiv*. <https://doi.org/10.48550/arXiv.2410.07959>
- [3] Novelli, C., Casolari, F., Rotolo, A., Taddeo, M., & Floridi, L. (2024). AI risk assessment: A scenario-based, proportional methodology for the AI Act. *Digital Society*, 3(1), 13. <https://doi.org/10.1007/s44206-024-00095-1>
- [4] McLachlan, S., & Webley, L. C. (2021). Visualisation of law and legal process: An opportunity missed. *Information Visualization*, 20(2–3), 192–204. <https://doi.org/10.1177/14738716211012608>

- [5] Gotsch, C., & Puchan, J. (2024). Harmonising innovation and governance: A lifecycle model for high-risk AI systems under the European AI Act. In *AKWI Jahrestagung 2024* (pp. 79–95). <https://doi.org/10.18420/AKWI2024-006>
- [6] Costa, M. Z., Guizzardi, G., & Almeida, J. P. A. (2022). On capturing legal knowledge in ontology and process models combined. In *35th International Conference on Legal Knowledge and Information Systems, JURIX 2022*, 267–272. <https://doi.org/10.3233/FAIA220478>
- [7] Waddington, M. (2021). Rules as code. *Law Context: A Socio-Legal J*, 37(1), 179–186. <https://doi.org/10.26826/law-in-context.v37i1.134>
- [8] Huggins, A., Burdon, M., Witt, A., & Suzor, N. (2022). Digitising legislation: Connecting regulatory mind-sets and constitutional values. *Law, Innovation and Technology*, 14(2), 325–354. <https://doi.org/10.1080/17579961.2022.2113670>
- [9] Mohun, J., & Roberts, A. (2020). *Cracking the code: Rule-making for humans and machines. OECD Working Papers on Public Governance* (Report No.42). OECD Publishing. <https://doi.org/10.1787/3afe6ba5-en>
- [10] Mowbray, A., Chung, P., & Greenleaf, G. (2023). Representing legislative rules as code: Reducing the problems of ‘scaling up’. *Computer Law & Security Review*, 48, 105772. <https://doi.org/10.1016/j.clsr.2022.105772>
- [11] Liu, H., Gegov, A., & Cocea, M. (2016). Rule-based systems: A granular computing perspective. *Granular Computing*, 1(4), 259–274. <https://doi.org/10.1007/s41066-016-0021-6>
- [12] Taylor, J., Fish, A., Vanthienen, J., & Vincent, P. (2013). *Emerging Standards in Decision Modeling — an Introduction to Decision Model & Notation*. Semantic Scholar. <https://api.semanticscholar.org/CorpusID:203597161>
- [13] Di Martino, B., Graziano, M., & Colucci Cante, L. (2024). Semantic, business process and natural language processing for ebuilding. In *International Conference on Complex, Intelligent, and Software Intensive Systems* 373–382. https://doi.org/10.1007/978-3-031-70011-8_35
- [14] Ghlala, R., Kodja, Z., & Ben Said, L. (2022). Enhancing decision-making consistency in business process using a rule-based approach: Case of business intelligence process. *Journal of Telecommunications and the Digital Economy*, 10(2), 44–61. <https://search.informit.org/doi/10.3316/informit.587137889753458>
- [15] Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1(1), 81–106. <https://doi.org/10.1007/BF00116251>
- [16] Mingers, J. (1989). An empirical comparison of pruning methods for decision tree induction. *Machine learning*, 4(2), 227–243. <https://doi.org/10.1023/A:1022604100933>
- [17] Jennath, H. S., & Asharaf, S. (2025). Explainable Optimal Random Forest model with conversational interface. *Engineering Applications of Artificial Intelligence*, 145, 110134. <https://doi.org/10.1016/j.engappai.2025.110134>
- [18] Balcan, M. F., & Sharma, D. (2024). Learning accurate and interpretable decision trees. In *The 40th Conference on Uncertainty in Artificial Intelligence*, 288–307. <https://openreview.net/forum?id=skdlnUYRzQ>
- [19] Gruber, T. R. (1995). Toward principles for the design of ontologies used for knowledge sharing? *International journal of human-computer studies*, 43(5–6), 907–928. <https://doi.org/10.1006/ijhc.1995.1081>
- [20] Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge acquisition*, 5(2), 199–220. <https://doi.org/10.1006/knac.1993.1008>
- [21] Staab, S., & Studer, R. (2009). *Handbook on Ontologies* (2nd ed.). Germany: Springer Berlin. <https://doi.org/10.1007/978-3-540-92673-3>
- [22] Patel, A., & Debnath, N. C. (2024). A comprehensive overview of ontology: Fundamental and research directions. *Current Materials Science: Formerly as Recent Patents on Materials Science*, 17(1), 2–20. <https://doi.org/10.2174/2666145415666220914114301>
- [23] Bibi, N., Rana, T. A., Maqbool, A., Mirza, A., Iqbal, Z., Khan, M. A., . . . , & Damaševičius, R. (2023). Web semantics and ontologies-based framework for software component selection from online repositories. *International Journal of Web and Grid Services*, 19(3), 318–349. <https://doi.org/10.1504/IJWGS.2023.133503>
- [24] Eiter, T., Ianni, G., Polleres, A., Schindlauer, R., & Tompits, H. (2006). Reasoning with rules and ontologies. In P. Barahona, F. Bry, E. Franconi, N. Henze, & Sattler. U (Eds.), *Reasoning Web International Summer School 2006* (pp. 93–127). Springer Berlin Heidelberg. https://doi.org/10.1007/11837787_4
- [25] Stănescu, G., & Oprea, S. V. (2025). Recent trends and insights in semantic web and ontology-driven knowledge representation across disciplines using topic modeling. *Electronics*, 14(7), 1313. <https://doi.org/10.3390/electronics14071313>
- [26] Governatori, G., Rotolo, A., Sartor, G., Gabbay, D., Horty, J., Parent, X., . . . , & van der Torre, L. (2021). Logic and the law: Philosophical foundations, deontics, and defeasible reasoning. *Handbook of Deontic Logic and Normative Reasoning*, 2, 655–760.
- [27] Gebser, M., Kaminski, R., Kaufmann, B., & Schaub, T. (2013). *Answer Set Solving in Practice*. USA: Morgan & Claypool Publishers. <https://doi.org/10.1007/978-3-031-01561-8>
- [28] Hanif, H., Constantino, J., Sekwenz, M. T., Van Eeten, M., Ubacht, J., Wagner, B., & Zhauniarovich, Y. (2024). Navigating the EU AI act maze using a decision-tree approach. *ACM Journal on Responsible Computing*, 1(3), 1–16. <https://doi.org/10.1145/3677174>
- [29] Golpayegani, D., Pandit, H. J., & Lewis, D. (2022). Airo: An ontology for representing AI risks based on the proposed EU AI Act and iso risk management standards. In *Towards a Knowledge-Aware AI: SEMANTiCS 2022—Proceedings of the 18th International Conference on Semantic Systems*, 51–65. <https://doi.org/10.3233/SSW220008>
- [30] Hernandez, J., Golpayegani, D., & Lewis, D. (2025). An open knowledge graph-based approach for mapping concepts and requirements between the EU AI Act and international standards. *AI and Ethics*, 5(5), 4463–4474. <https://doi.org/10.1007/s43681-025-00708-6>
- [31] Lawniczak, L., & Benz Müller, C. (2025). Logical modalities within the European AI Act: An analysis. In *Proceedings of the Twentieth International Conference on Artificial Intelligence and Law*, 348–353. <https://doi.org/10.1145/3769126.3769202>
- [32] Francesconi, E., & Governatori, G. (2023). Patterns for legal compliance checking in a decidable framework of linked open data: E. Governatori. *Artificial Intelligence and Law*, 31(3), 445–464. <https://doi.org/10.1007/s10506-022-09317-8>
- [33] Gallese, C., & Falletti, E. (2025). Dataset assessment tools for the AI Act compliance of high-risk AI Systems. *Journal of Law, Market & Innovation*, 4(3), 653–682. <https://doi.org/10.13135/2785-7867/12840>
- [34] Dumas, M., La Rosa, M., Mendling, J., & Reijers, H. A. (2018). Introduction to business process management. In M. Dumas, M. La. Rosa, J. Mendling, & H. A. Reijers

- (Eds.), *Fundamentals of Business Process Management* (pp. 1–33). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-56509-4_1
- [35] Dumas, M., Fournier, F., Limonad, L., Marrella, A., Montali, M., Rehse, J. R., . . . , & Weber, I. (2023). AI-augmented business process management systems: a research manifesto. *ACM Transactions on Management Information Systems*, 14(1), 1–19. <https://doi.org/10.1145/3576047>
- [36] Huy, P. Q., & Phuc, V. K. (2025). Unveiling how business process management capabilities foster dynamic decision-making for effectiveness of sustainable digital transformation. *Business Process Management Journal*, 31(8), 67–103. <https://doi.org/10.1108/BPMJ-06-2024-0467>
- [37] Chapela-Campa, D., Benchechrout, I., Baron, O., Dumas, M., Krass, D., & Senderovich, A. (2025). A framework for measuring the quality of business process simulation models. *Information Systems*, 127, 102447. <https://doi.org/10.1016/j.is.2024.102447>
- [38] Governatori, G. (2024). An ASP implementation of defeasible deontic logic. *KI-Künstliche Intelligenz*, 38(1), 79–88. <https://doi.org/10.1007/s13218-024-00854-9>
- [39] Huggins, A., Witt, A., & Burdon, M. (2024). Digital distortions arising when encoding regulation. *Research Features*, 151, 1–3. <https://doi.org/10.26904/RF-151-6106305305>
- [40] Quaranta, M., Amantea, I., Governatori, G., Molinari, M., Billi, M., Galli, F., & Rotolo, A. (2025). *AI Compliance frameworks: A handbook for compliance tools based on AI Acts' requirements*. SSRN. <https://dx.doi.org/10.2139/ssrn.5555260>
- [41] Savvadelli, E., Kiouvrekis, Y., & Kokkinaki, A. (2025). A literature review on rule-based systems as decision support systems. In *International Symposium on Human Aspects of Information Security and Assurance*, 376–388. https://doi.org/10.1007/978-3-032-02504-3_26
- [42] Liu, X., Huang, H., Li, Y., Xia, Z., & Zhao, S. (2026). Information-Driven Rule Reduction in Belief Rule Bases for Complex System Modeling. *Entropy*, 28(5), 578. <https://doi.org/10.3390/e28050578>
- [43] Governatori, G. (2026). Weak permission is not captured by well-founded, grounded, and stable semantics in the presence of deontic conflicts. In *Proceedings of the International Conference on Artificial Intelligence and Law (ICAIL 2026)*.

How to Cite: Quaranta, M., Bucci, A., & Amantea, I. A. (2026). Assessing the AI Act from the Root to the Leaf: A Decision Tree Methodology. *Journal of Computational Law and Legal Technology*. <https://doi.org/10.47852/bonviewJCLLT620210181>