A GSN-Based Requirement Analysis of the EU AI Regulation

Natsuki Hayama[1], Yoriyuki Yamagata[2], Hideaki Nishihara[3], and Yutaka Matsuno✉[1]

[1] Department of Computer Science, Graduate School of Science and Technology, Nihon University, JAPAN
csna24011@g.nihon-u.ac.jp, matsuno.yutaka@nihon-u.ac.jp
[2] Graduate School of Engineering, University of Fukui, JAPAN
yoriyuki@u-fukui.ac.jp
[3] National Institute of Advanced Industrial Science and Technology (AIST), Cyber Physical Security Research Institute, JAPAN
h.nishihara@aist.go.jp

Abstract. Artificial intelligence (AI) is now embedded in many domains, yet persistent concerns remain about opaque decision processes, training-data bias, and impacts on privacy and human rights. The European Union's AI Act addresses these risks through a risk-based framework that imposes stringent obligations on systems classified as high-risk. The Act's voluminous and intricate provisions make it difficult for engineers and compliance teams to obtain an integrated view from the legal text alone. This paper employs Goal Structuring Notation (GSN) to translate the AI Act's high-risk requirements into a layered argument model. The resulting GSN diagram exposes the logical links between individual articles, annexes, and the supporting evidence needed for conformity. The approach improves traceability, highlights documentation gaps, and offers a maintainable structure for tracking future amendments to the legislation. The resulting GSN diagram not only exposes the logical links between individual articles but also serves as a common ground for stakeholders to discuss and negotiate AI Act compliance.

Keywords: Assurance Cases · Goal Structuring Notation (GSN) · EU AI Act

1 Introduction

Artificial intelligence (AI) now underpins critical functions in healthcare, finance, mobility, education, and public safety. Its rapid growth, however, has amplified concerns about opaque decision logic, data bias, and risks to privacy and human rights. These issues have triggered calls for verifiable safety, reliability, and fairness across the AI lifecycle.

The European Union's Artificial Intelligence Act (EU AI Act) [7] answers this need with a risk-based framework that imposes stringent obligations—such as risk management, data governance, and conformity assessment—on systems

flagged as high-risk. Because the Act's provisions are both voluminous and intricate, stakeholders struggle to see how individual clauses connect or what concrete evidence is required for compliance.

We use Goal Structuring Notation (GSN) [11] to organise the AI Act's high-risk requirements into a layered argument. This structure explicitly links each regulatory goal to its supporting strategies and evidence, providing a clear overview of dependencies across articles and annexes. This study contributes a comprehensive GSN model that integrates mandatory governance, risk management, and post-market monitoring requirements of the AI Act into a unified assurance argument. This paper assumes that providers will be the primary users of the GSN model. Therefore, the model incorporates providers' obligations as nodes, enabling clear visualization of their compliance duties. The whole GSN diagram can be seen in [9], using D-Case Communicator, a web-based GSN editor [13]. † While this GSN model represents one interpretation of the EU AI Act requirements, it aims to facilitate constructive discussions among various stakeholders—including engineers, compliance teams, and regulators—by providing a shared visual framework.

This paper is structured as follows. Section 2 introduces the EU Artificial Intelligence Act. Section 3 positions this research in the context of related work. Subsequently, Section 4 describes the methodology employed in this study, detailing the steps involved in translating the provisions of the EU AI Act into GSN. In Section 5, we present the GSN model, illustrating how the requirements for risk management and conformity assessment of high-risk AI can be modeled. Section 6 examines the usefulness and limitations of the GSN model, as well as prospects for future development. Finally, Section 7 concludes the paper.

2 The EU Artificial Intelligence Act

In response to the societal impacts and risks posed by AI technology, the European Union (EU) has enacted the world's first comprehensive AI legislation, known as the Artificial Intelligence Act. Under this regulation, AI systems are classified according to four risk categories—prohibited, high-risk, limited-risk, and minimal-risk—with stringent requirements such as risk management and data governance imposed on those deemed high-risk. A key feature is that high-risk AI systems intended for market release must undergo conformity assessment procedures, typically involving CE marking, thereby clarifying the provider's responsibility to ensure safety from development through post-deployment. As the Act allows for substantial fines—similar to the General Data Protection Regulation (GDPR)—it is expected to have a significant global impact, potentially influencing international AI governance standards. In this paper, we focus on the high-risk AI requirements of the EU AI Act, which center on a risk-based approach and obligations placed on providers, and examine how Goal Structuring Notation can be used to visualize these requirements.

† https://x.gd/laKX6

3 Related Work

Goal Structuring Notation (GSN) [6] is a graphical notation for assurance cases that logically demonstrate properties—such as safety—in safety-critical systems by visualizing the relationship between "goals" and the supporting "evidence (Solution)." Its core elements comprise a top-level claim ("Goal"), the "Strategy" that explains how the goal is decomposed or justified, the "Context" that provides premises for goals and strategies, and the "Solution" that supplies the evidence. Typical practice applies a Strategy to decompose a Goal into multiple sub-goals that all must be satisfied (AND-decomposition). For example, when claiming that "the system is safe", the goal may be divided into sub-goals such as hardware safety, software safety, and operational safety procedures, each supported by documents or test results. Where fulfillment of any of several conditions suffices, we adopt an OR-decomposition, following the pattern construct proposed by [5]. Recent extensions target AI and data-driven systems: BIG Argument framework is introduced for AI safety cases, demonstrating how GSN patterns can be applied to machine-learning pipelines [8].

Recently, a variety of AI-governance instruments have emerged in addition to the EU AI Act. These include the process-oriented ISO/IEC 42001 [1], the ISO/IEC JTC 1/SC 42 family of technical specifications [3], the NIST AI Risk Management Framework [15], the Singapore Consensus [4], and the Council of Europe Convention on AI [2]. Among them, the EU AI Act and the Council of Europe Convention are legally binding regulations, making compliance essential for AI-system providers.

Fraunhofer IKS's white paper [10] surveys the EU AI Act's risk-based framework, analyses the high-risk obligations in Articles 9–15, and maps them to an extended ISO/IEC 25059 quality-attribute model. It introduces a contract-based design approach that allocates responsibilities along the AI value chain and, through three sector-specific case studies, demonstrates how local GSN fragments can justify attributes such as traceability, fairness, and explainability. In contrast, our work concentrates on the provider perspective and translates the entire set of high-risk provisions—covering the classification logic of Article 6 and Annex III as well as the operational duties in Articles 16–21 and 43–49—into a single, six-layer GSN model. To the best of our knowledge, this is the first end-to-end mapping that unifies mandatory governance, risk-management, and post-market monitoring requirements within one coherent assurance argument.

Myklebust et al.'s The AI Act and The Agile Safety Plan [14] provides a practical framework for integrating the AI Act's high-risk requirements with agile development processes. The book maps traditional safety case structures to Agile Safety Plans, clarifies stakeholder responsibilities (Articles 16-25), and aligns technical requirements (Articles 9-15) with existing standards such as ISO/IEC 42001 and NIST AI RMF. Their emphasis on maintaining a "living document" through iterative updates complements our GSN-based approach by addressing the operational aspects of continuous compliance management.

4 Proposed Method

This section outlines the method for modeling and analyzing the main requirements of the EU AI Act using GSN. We review the sections of the Act and Annex III relevant to high-risk AI, which form the focal point of the paper, and then introduce our approach.

We adapt GSN models to legal analysis of large-scale, complex regulatory frameworks. In the EU AI Act, where multiple provisions and annexes comprehensively define requirements, and where the classification of AI systems and conformity assessment procedures vary according to different risk levels, GSN-based visualization is expected to potentially facilitate understanding and simplify maintenance efforts.

We propose a procedure for using GSN to model high-risk AI requirements and provider obligations under the EU AI Act as follows. These steps are based on [6].

1. **Extract and categorize the AI Act's requirements**
   Review the contents of the AI Act that deal with the management of high-risk AI and the annexes listing high-risk AI systems and organize the requirements.
2. **Design the basic structure of the GSN model**
   Define a top-level goal, and model each requirement or system domain as a sub-goal. If the Act introduces conditional steps or additional processes, decompose them into sub-goals and link them with AND/OR strategies to capture the branching logic. When decomposing a higher-level goal, we invariably insert a Strategy node before placing the corresponding sub-goals beneath it. Thus each branch follows the sequence Goal → Strategy → Sub-Goal, which makes the rationale for the decomposition explicit and allows readers to grasp the role of every sub-goal at a glance.
3. **Assign evidence, and assumptions**
   Link the relevant audit reports and internal documents as evidence, and cite the text of the EU AI Act and related guidelines as context or assumptions.

4.1 Provisions of the EU AI Act Incorporated into the GSN Model

Proposed by the European Commission (EC) and adopted in March 2024, the EU AI Act [7], classifies AI systems into four risk categories: prohibited, high-risk, limited-risk, and minimal-risk, and imposes especially stringent requirements on those deemed high-risk. In this paper, we focus on the regulatory requirements for high-risk AI under the AI Act, addressing the stringent regulatory content involved. We construct a GSN model that incorporates the relevant provisions on provider obligations, enabling providers developing high-risk AI systems to verify that they meet these regulatory requirements. By taking this approach, we aim to ensure that providers can carry out the necessary checks to maintain regulatory compliance.

First, Articles 9 to 15 define the main elements of high-risk AI requirements, including risk management (Article 9), data governance (Article 10), technical documentation (Article 11), record-keeping (Article 12), transparency (Article 13), human oversight (Article 14), and accuracy and robustness (Article 15). These articles set forth the specific standards that high-risk AI systems must meet. Next, Articles 16 to 21, along with Articles 43, 47, 48, and 49, detail the obligations of providers, including information disclosure, monitoring, the implementation of corrective measures, conformity assessment procedures such as CE marking, and declarations of compliance.

Additionally, Annex III supplements Article 6 by listing high-risk domains such as employment, finance, healthcare, public services, and law enforcement, thus clarifying which AI systems qualify as high-risk. These sections of the Act are especially significant for companies and developers striving for compliance, making them essential sub-goals in the GSN model proposed in this study.

Because the relevant provisions are extensive, we focus primarily on the articles mentioned above and does not consider references made to other legislation within them. As the EU AI Act may undergo further amendments, we base our discussion on the version adopted as of March 2024.

5 Structuring the EU AI Act with GSN

In this section, we introduce our GSN model. The whole GSN diagram can be seen in [9], using D-Case Communicator [13] [†] .

We divide the constructed GSN into the following three perspectives and explain them in sequence:

1. **Setting the Top Goal**
2. **Determining high-risk AI eligibility**
3. **Proper Management of High-Risk AI**

5.1 Setting the Top Goal

Fig. 1 shows the top-level goal and its sub-goals. At the highest level, we define:

G1: "X is compliant with the AI Act as a high-risk AI."

This top-level goal indicates whether the target system (X) meets the legal requirements and ensures its requirements in development and operation.

Below G1, we introduce a strategy node stating:"Separate the discussion into two parts: classification of the target AI system (X) and how it is managed." allowing us to determine which risk category the system falls into and derive the subsequent branching.

To achieve the top-level goal, we split the argument into two sub-goals:

G2: "X is a high-risk AI."

[†] https://x.gd/laKX6

This sub-goal argues that the target AI system fits the definition of a high-risk AI system under the EU AI Act, thereby clarifying that relevant regulatory requirements apply.

G3: "X is properly managed as a high-risk AI."

This sub-goal argues that the high-risk AI system is being managed in accordance with the requirements, thereby ensuring its reliability and safety.

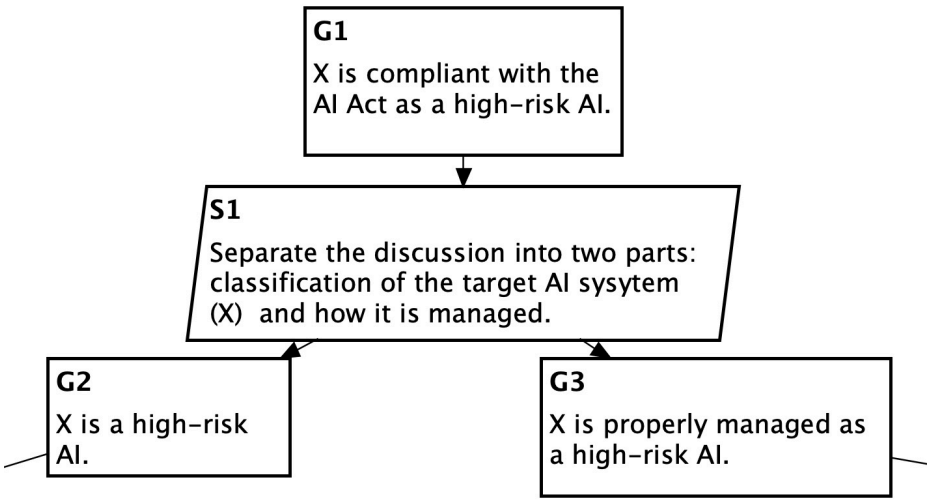

Fig. 1. Top level of the GSN model

5.2 Determining high-risk AI eligibility

We show the sub tree of the GSN model G2: "X is a high-risk AI." in Fig.2. The strategy node states: Analyze the system's functions, applications, and risk level to demonstrate that X is a high-risk AI under the AI Act.

Continuing, we set sub-goals to establish that the AI is not prohibited and that it conforms to the categories indicated by the AI Act. Then, after G4, we verify whether it falls under Annex III or whether it is subject to the evaluation obligations for products covered by Annex I, thereby clearly demonstrating that X is classified as a high-risk AI.

Here, we use an OR branch to represent the fact that if the system appears on either the list of EU harmonization legislation in Annex I or within the domains listed in Annex III, it is deemed high-risk AI.

As shown in Fig.2, under Article 6 of the EU AI Act, high-risk AI systems are classified according to Annexes I and III. Therefore, we set the classification requirements in these annexes as sub-goals and also include a sub-goal showing

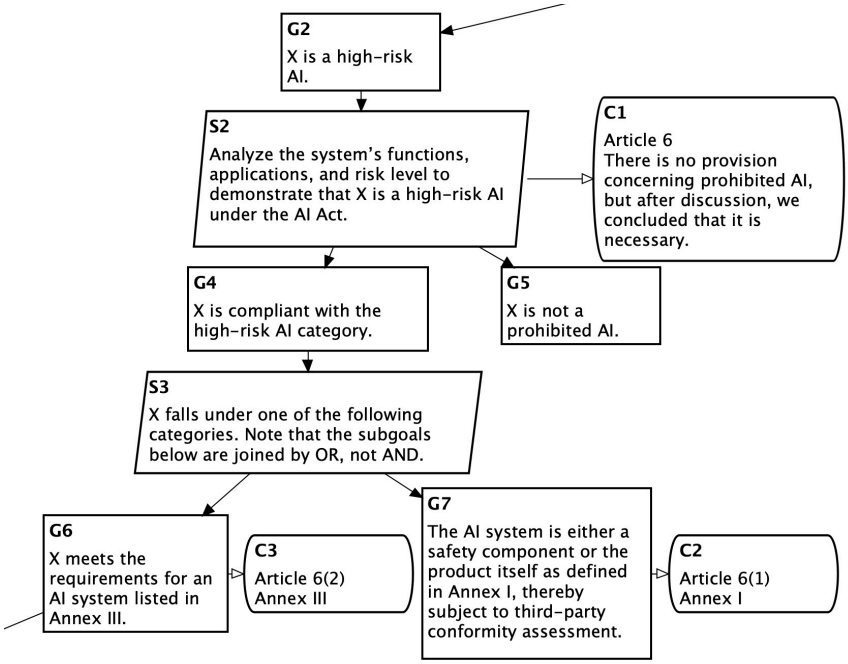

Fig. 2. GSN part for arguing high-risk AI system

that "X is not a prohibited AI." Because Annex I references other legislation outside the scope of this paper, we do not go into detail here.

Fig. 3 illustrates part of the GSN model for Annex III.

If a system falls under any one of the domains specified in Annex III (such as biometrics, critical infrastructure, education and vocational training), it can be judged high-risk AI. Hence, we use an OR branch.

G6: "X meets the requirements for an AI system listed in Annex III." The accompanying strategy node is S4: "X falls under one of the domains indicated below. " enabling branching based on the domain in which the AI operates. In Fig. 3, we show two sub-goals:

G8: "It is an AI system used in the field of biometrics."

G9: "It is an AI system used as a safety component in the management and operation of critical infrastructure."

If the relevant article further subdivides those domains by use case, we can add new strategy nodes and sub-goals under G8, as shown in Fig. 3.

By modeling the classification process in the GSN model, providers can verify the domain or specific purpose under which their AI system falls, thereby making it clear whether the AI system needs to comply with high-risk AI requirements.

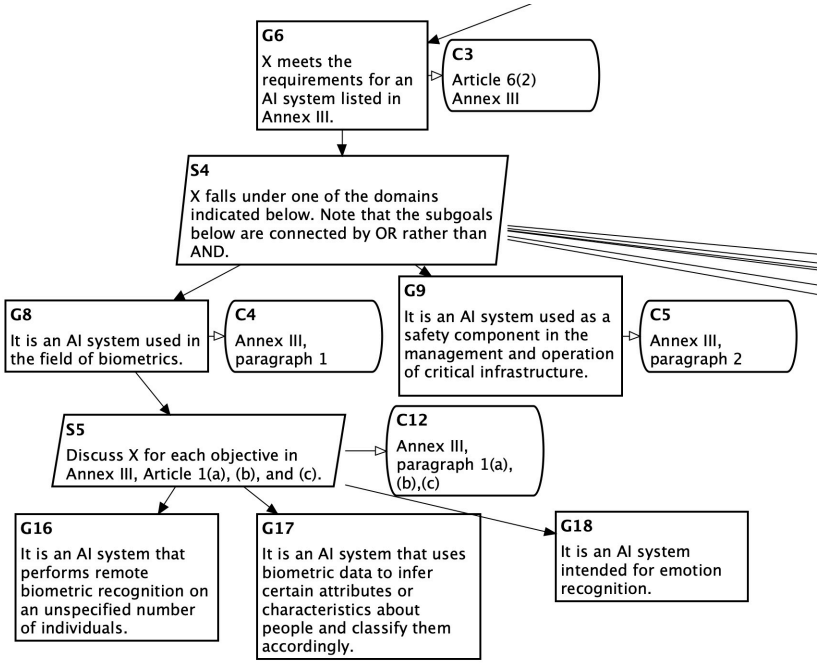

Fig. 3. GSN part for Annex III requirements

5.3 Proper Management of High-Risk AI

This subsection explains the part of the GSN model related to G3: "X is properly managed as a high-risk AI" (Fig. 1), with the correponding GSN subtree shown in Fig. 4.

Under G3, we introduce a strategy node:

S12: "Divide high-risk AI into three parts—procedures, compliance with requirements, and operation/monitoring—and discuss each separately."

This allows us to visualize the entire lifecycle of high-risk AI step by step, logically organizing the requirements and evidence for each phase. As a result, it becomes clearer which documentation and audit procedures are required at each phase.

Specifically, we decompose G3 into three sub-goals:

G40: "Perform documentation and procedures for the high-risk AI system."

G41: "Comply with the requirements that high-risk AI systems must fulfill."

G42: "Continuously operate and monitor the high-risk AI system."

Each of these sub-goals addresses certain technical requirements from Articles 9–15, along with the provider obligations. Note that these three categories are not explicitly defined in the EU AI Act; rather, we introduced them to improve the model's overall readability. If we were to break down every single requirement into its own sub-goal, the model could become unwieldy. Instead, grouping by

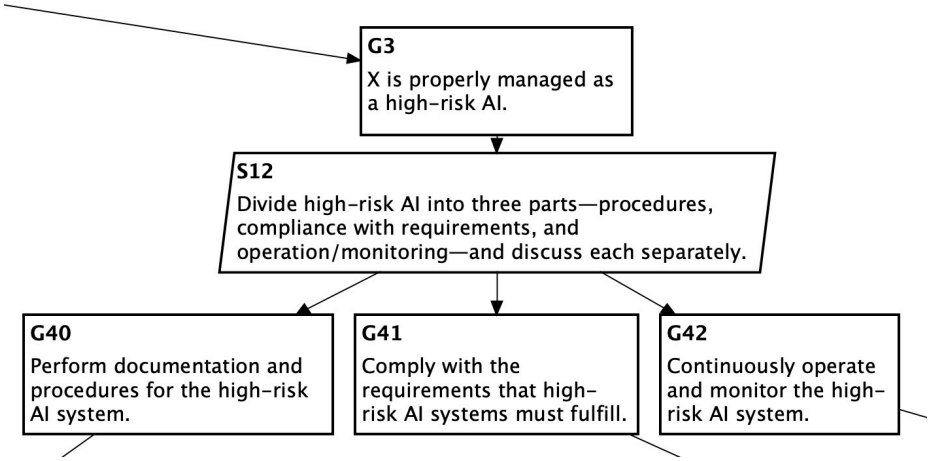

Fig. 4. GSN part for managing the risks of high-risk AI systems

"documentation and procedures," "standards compliance," and "operation and monitoring" balances readability and comprehensiveness, while also offering an intuitive representation of the risk management process.

Next, Fig. 5 focuses on the G40 sub-goal, which organizes the documentation and procedural requirements for high-risk AI.

For G40, we place a strategy node:

S13: "Ensure legal compliance and prepare the system for market entry."

This clarifies the process of meeting the legal requirements before launching the service, helping to ensure that the high-risk AI can enter the market smoothly. Following S13, we include seven sub-goals, of which Figure 5 shows two as examples:

G44: "In accordance with Article 48, affix a CE mark to the AI system demonstrating compliance with this regulation; if that is not possible, affix it to the packaging or accompanying documentation."

G45: "Complete an appropriate conformity assessment procedure before placing the AI system on the market or offering the service."

In the case of G45, the text includes additional branches depending on the domain in which the AI system is used, so the GSN continues beneath G45 with further Strategy and sub-goal nodes.

By hierarchically displaying the requirements at each stage in the GSN model, it becomes much easier to visually comprehend the AI Act's provisions—rather than relying solely on text.

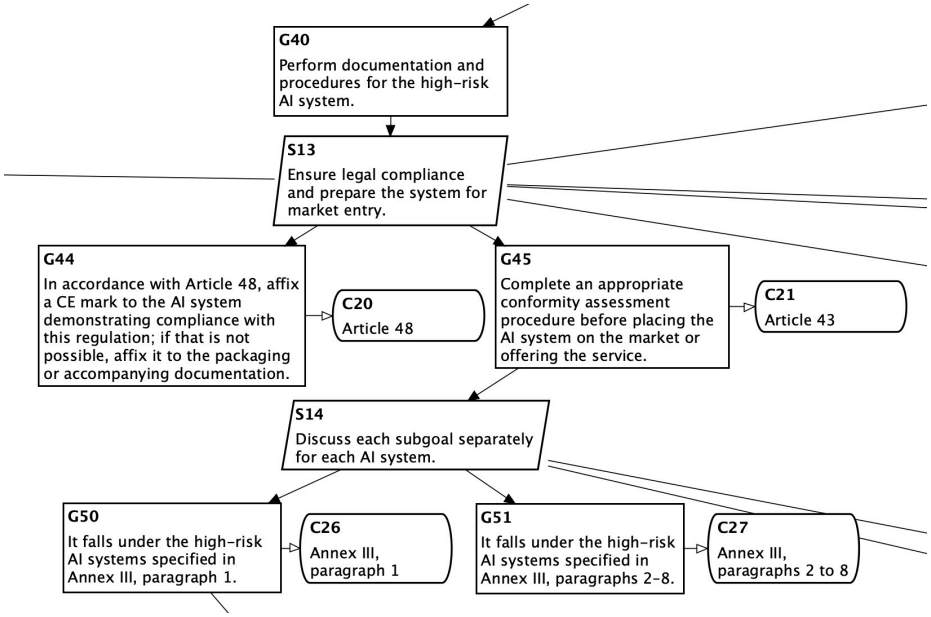

Fig. 5. GSN part for documentation and procedures

6 Discussion

6.1 GSN clarifies EU AI Act

In this paper, after organizing the relevant articles of the EU AI Act, we initially focused on 20 articles. Moreover, if each clause up to Article 13(3)(b)(i) is fully subdivided, the total can reach up to 156 items. On the other hand, our GSN model contains 129 nodes in total, with 50 of those being sub-goals at the lowest level.

The difference in counts is due to the fact that we proactively decomposed clauses that could be readily isolated, so the number of sub-goals exceeds the article count (20). Conversely, because certain sections were left undisaggregated, the total remains far below the theoretical maximum (156). This approach preserves the necessary level of specificity while preventing the model from becoming excessively large.

Two factors contribute to these comprehensive structures. The first is that in some instances, individual clauses, such as Article 13(3)(b)(i), were not subdivided in detail but rather consolidated into a single sub-goal. While this approach raises the level of abstraction and simplifies the discussion, practitioners seeking a more detailed understanding may need to refer back to the original text. The second factor is the more extensive subdivision of specific articles, such as those in Annex III. In these cases, multiple elements within a single article are treated as distinct conditions or domains, which can prompt the creation of additional sub-goals in the GSN model to capture these details.

As a result, some articles have been broken down into multiple sub-goals, whereas certain other articles remain relatively abstract. However, creating a strict one-to-one mapping of every clause would likely make the model unwieldy. Thus, the decision of how to "group" or consolidate requirements inevitably depends to some extent on the discretion of the model developer. For every clause-related sub-goal, the referenced article number is recorded in the corresponding Context node, allowing the original text to be traced at any time. This mechanism prevents information loss caused by abstraction.

6.2 Depth of the GSN Model

In this study, the final GSN model has six levels of hierarchy. It takes six layers to move from the top-level goal down to the lowest-level sub-goals because of the numerous and varied requirements in the AI Act. While a deeper hierarchy allows for more granular representation of legal requirements, it also makes the diagram more complex and harder to view in its entirety. A depth of six levels is relatively substantial and demands a certain learning cost for readers. It also necessitates careful attention when the Act is amended or the system is updated, as it is crucial to identify exactly which layer needs revision.

In summary, having six levels of depth does help capture the complexity of the AI Act but also highlights issues of model readability and maintainability.

6.3 Challenges and limitations of the GSN Model

The following challenges remain for modeling the EU AI Act with GSN.

Maintenance burden Even after its final enforcement, the EU AI Act may be subject to revisions or additional guidelines. Any GSN model created from it must be continually updated to reflect such changes in law and internal processes. This is particularly important in places where OR branching or complex subdivisions are used, as even a small amendment to one article could have ramifications throughout the entire model. A system to keep the GSN model up to date is therefore essential.

Loss of readability due to model expansion This paper focuses primarily on high-risk AI provisions, along with provider obligations and conformity assessment procedures. By carefully limiting the number of nodes in our GSN model, we have managed to simplify its structure and enhance its overall usability. However, the model spans six hierarchical levels, which could make future maintenance and updates increasingly cumbersome. Moreover, if we were to include requirements for other stakeholders—such as deployers or importers—or other risk categories, the model would become large and potentially less readable. To address this complexity, model designers could separate GSN diagrams by risk level, or deliberately consolidate and modularize requirements, thus preventing the model from becoming unwieldy.

Limitations of GSN and the need for extensions Although we used GSN here to visualize the EU AI Act, GSN is not a strictly formal logic system; it is a notation for illustrating "how" a system is supposed to achieve or demonstrate certain objectives or properties.

For example, consider an AI system that uses biometric data in real time to authenticate and monitor examinees taking an online test remotely. Such a system may fall under both the "biometrics" and "education" domains specified in Annex III. However, if we use an exclusive OR branch ("either/or") in our S3 node, we cannot effectively capture the fact that the system actually falls under both domains simultaneously. Handling conditions that might apply in multiple ways (e.g., "if...then" for more than one option) would require enhancements to GSN, allowing for logical relationships in which multiple requirements can be met concurrently.

While our GSN model effectively visualizes EU AI Act requirements, it is important to note that GSN visualization alone does not guarantee actual compliance or system safety, as highlighted by Leveson's critique of safety assurance methods [12].

6.4 GSN as a discussion facilitator

This GSN model serves not merely as a compliance checklist but as a platform for dialogue. By making implicit relationships explicit and providing a common visual language, it enables stakeholders with different backgrounds to engage in meaningful discussions about EU AI Act implementation. The model's incompleteness, rather than being a weakness, invites collaborative refinement and adaptation to specific contexts.

7 Conclusion and Future Work

In this paper, we proposed a method for logically and hierarchically visualizing the complex requirements set forth in the EU AI Act—particularly those concerning high-risk AI—by applying the Goal Structuring Notation (GSN). Our main contribution lies in using GSN to represent the diverse legal requirements, which have traditionally been difficult to grasp from text alone, thereby making it easier to comprehensively understand both the interrelationships among these requirements and the overall flow of the risk management process. Importantly, this GSN model is offered as one possible interpretation of the EU AI Act, intended to stimulate discussion and collaborative understanding rather than prescribe a definitive compliance approach.

At the same time, several challenges remain, including adapting to legislative changes or additional guidelines, as well as maintaining readability when dealing with extensive or complex provisions. As a future prospect, we believe that extending the logical foundation of GSN will be beneficial. For example, by introducing mechanisms to represent OR branching and relationships that satisfy multiple conditions simultaneously, the GSN proposed in this study can be

further developed. Formalizations and implementation of such GSN extensions will reduce the cost of compliance with EU AI Act.

Acknowledgments. This research was supported by JSPS KAKENHI Grant Number JP23H03376 (2023).

References

1. ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system (2023)
2. Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (2024)
3. ISO/IEC JTC 1/SC 42 — Artificial intelligence (Technical committee homepage) (2025)
4. The Singapore Consensus on Global AI Safety Research Priorities (2025)
5. Alexander, R., Kelly, T., Kurd, Z., McDermid, J.: Safety cases for advanced control software: Safety case patterns. Tech. rep., Department of Computer Science, University of York (2007)
6. Assurance Case Working Group: Goal structuring notation community standard version 3 (May 2021), downloaded from https://scsc.uk/r141C:1?t=1
7. European Union: Regulation (eu) 2024/1689 of the european parliament and of the council of 13 june 2024 laying down harmonised rules on artificial intelligence (artificial intelligence act) 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 (2024), accessed: 2025-04-28
8. Habli, I., Hawkins, R., Paterson, C., Ryan, P., Jia, Y., Sujan, M., McDermid, J.: The big argument for ai safety cases. arXiv preprint **arXiv:2503.11705** (2025)
9. Hayama, N.: A GSN Model for EU AI regulation, https://x.gd/laKX6
10. Heidemann, L., Herd, B., Kelly, J., Mata, N., Tsai, W., Zafar, S., Zamanian, A.: The European Artificial Intelligence Act Overview and Recommendations for Compliance (2024)
11. Kelly, T., Weaver, R.: The Goal Structuring Notation - a safety argument notation. In: Proc. of the Dependable Systems and Networks 2004, Workshop on Assurance Cases (2004)
12. Leveson, N.: The use of safety cases in certification and regulation. In: ESD Working Paper Series, Boston: MIT (2011)
13. Matsuno, Y.: D-Case Communicator: A web based GSN editor for multiple stakeholders. In: Proc. ASSURE. LNCS, vol. 10489, pp. 64–69. Springer (2017)
14. Myklebust, T., Stålhane, T., Vatn, D.M.K.: The AI Act and The Agile Safety Plan (2025)
15. Tabassi, E.: Artificial intelligence risk management framework (ai rmf 1.0) (Jan 2023)