Complying with the EU AI Act

Jacintha Walters[(✉)] , Diptish Dey , Debarati Bhaumik ,
and Sophie Horsman

University of Applied Sciences Amsterdam, 1091 GC Amsterdam, The Netherlands
jacintha.walters@gmail.com

Abstract. The EU AI Act is the proposed EU legislation concerning AI systems. The goal of this paper is to determine on which areas organizations should focus with regards to compliance with the AIA. This paper identifies several categories of the AI Act. Based on this categorization, a questionnaire is developed that serves as a tool to offer insights by creating quantitative data. Analysis of the data shows various challenges for organizations in different compliance categories. The influence of organization characteristics, such as size and sector, is examined to determine the impact on compliance. The paper will also share qualitative data on which questions were prevalent among respondents, both on the content of the AI Act as on the application. The paper concludes by stating that there is still room for improvement in terms of compliance with the AIA and refers to a related project that examines a solution to help these organizations.

Keywords: EU AI Act · compliance preparedness · organization challenges

1 Introduction

The EU AI Act (in this paper abbreviated as AIA) is a proposed regulation (law) by the European Commission that aims to regulate the application of artificial intelligence in the European Union. The proposed regulation was published in 2021 and is currently under review by the European Parliament and the Council of the European Union. It defines which AI systems are categorized as high-risk and the rules applicable before a high-risk AI can be used [1]. At the time of writing, it is unknown when the AIA will become in effect [2].

One of the challenges in complying with the AIA is that AI systems are developed and maintained by a chain of actors, including software developers, data scientists, and engineers. The challenges for organizations are further complicated because of the interdisciplinary character of legal, technical, and domain-specific responsibilities. For an organization to comply, it must be able to interpret the contours implied by the act and translate this information into relevant requirements.

This paper identifies areas where organizations face challenges when considering current and future compliance with the AIA. The following steps are undertaken to identify these areas. Initially, categories of concern within the AIA are identified. Based on this categorization, a questionnaire is constructed to gather insights into how organizations handle the requirements associated with each category. The questionnaire is further refined through expert reflection, and trial runs to ensure its effectiveness.

2   Relevant Literature

Usman *et al.* observes that organizations can be subject to multiple regulations, which may lead to several challenges. First, there can be conflicting requirements. Second, some regulations are not well-defined, leaving the development team unsure how to implement them. After implementation, it can be challenging to verify that the software system meets all the requirements [3].

Research on privacy regulations showed that many small medium enterprises (SMEs) do not possess sufficient knowledge of the regulations to achieve compliance. Besides the risk of fines, compliance is essential to sustain if organizations want to supply services to other compliant organizations [4].

There can be a significant difference in compliance for different sectors. And one study performed in Malaysia concluded that Government-owned organizations generally demonstrate a lower degree of compliance than other organizations [5].

Most existing research on the AIA has a theoretical perspective, focusing mainly on the quality of the content of the AIA rather than on the application. One study concludes that the AIA is a good attempt but has several weaknesses. For instance, many parts are ambiguous, making it hard for organizations to define rules to self-assess against [6]. Another study concludes that the AIA is generally well-constructed but advises that the proposal should not rely so heavily on internal controls. External oversight is a necessity [7].

A notable research gap exists regarding the future compliance of organizations with the AIA and their level of preparedness. The existing literature covers two parts. First, compliance with existing regulations like GDPR. Second, critical analysis of the content of the AIA. There is a lack of insight into how organizations will navigate compliance with the AIA and the extent to which they are prepared. This paper aims to address this research gap by providing insights into the level of preparedness and the challenges organizations will face in complying with the AIA.

3   Methodology

3.1   Identifying Categories in the AIA

Figure 1 shows an overview of the relevant documentation for AIA compliance. The AIA focuses on subjects such as technical documentation, user communication and risks to human rights and discrimination. However, to ensure a

manageable questionnaire size, a decision was made to exclude certain subjects discussed in the AIA. The subjects of robustness, cybersecurity, logging, reporting, and audit preparedness were omitted from the questionnaire. This exclusion was primarily driven by the need to reduce the questionnaire's length, making it more feasible for potential respondents to complete. Although relevant to AI development, these subjects are broader and primarily associated with IT development.

Based on this selection, Fig. 2 shows a hierarchical breakdown of relevant key subject areas from the AIA used as a basis for the questionnaire. The breakdown in Fig. 2 is a result of highlighting key subject areas of the AIA and breaking them down into categories. This was done by focusing on the parts that are most relevant for organizations and summarizing the important information.
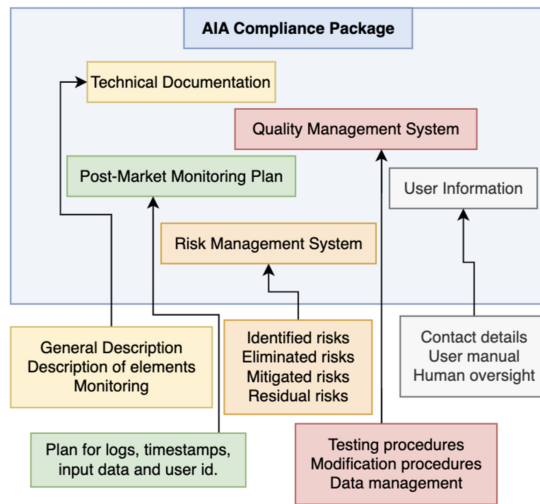

Fig. 1. Overview of the AIA compliance documentation

3.2 Creating and Refining the Questionnaire

The categories in Fig. 2 were used to create a questionnaire that assesses compliance with the AIA. The final questionnaire contains 5 parts: data & model internals, technical documentation, user communication, model monitoring (including human oversight), and risk management (including quality management and risk management).

The construction of the questionnaire involved several iterative steps to improve its validity and reliability. Existing questions were rephrased to transform open-ended questions into closed questions. Proxy questions were incorporated to ensure fair and reasonable responses. For example, the statement
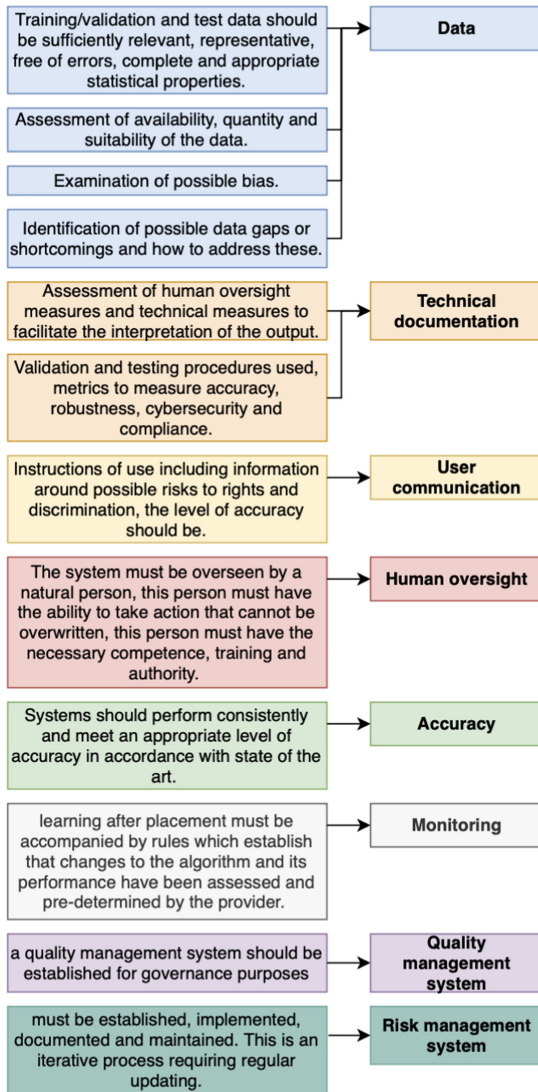
Fig. 2. Overview of key subject areas from the AIA

"My organization identifies and mitigates risks associated with a dataset" is supported by the question, "How often does your organization mitigate risks in a dataset?".

The questionnaire contains around 90 questions, which took respondents about 15 min to answer.[1] Feedback was gathered on the questionnaire from two organizations through an online interactive trial run. Most questions are 'state-

---

[1] Questionnaire: https://hva.eu.qualtrics.com/jfe/form/SV_9sFXWLoj5uFoaua.

ment questions', 'how often' questions and 'who' questions as can be seen in the template questions in Figs. 3, 4 and 5. "Statement questions" rely primarily on respondents' perspectives, whereas other questions are more objective. This observation is used to compute a "compliance score".

Q75

Professionals in my organization are aware that data can be biased.

Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree
○	○	○	○	○

Fig. 3. 'Statement...' questions

Q36

How often does your organization mitigate risks in a dataset?

○ In almost every project

○ Often

○ Seldom

○ This has never happened

Fig. 4. 'How often...' questions

Interactive interviews are conducted with multiple respondents based on the questionnaire. An online questionnaire is also circulated to obtain a broader range of responses. In total, seven responses were obtained through interactive interviews, supplemented by eight responses gathered online.

3.3 Response Rating

Each of the fifteen responses is rated using a three-point range. The scoring process aims to quantify the responses for each questionnaire category to enable numeric comparison. The questionnaire data is rated using a rule-based system. The rule-based system involves manually creating rules that are used to score each entry in the dataset automatically.[2]

The following categories from the questionnaire are used: data and model internals, technical documentation, user communication, model monitoring, and risk management. Generally, each question has a "perfect" answer worth 2 points, followed by "reasonable" answers worth 1 point. If a question had multiple options that should be selected, each option is worth 1/2 point. The remaining answers score 0 points. The perfect answer aligns closely with the requirements stated in the AI Act. The point distribution for each question is summarized in Fig. 6.

---

[2] Code: https://gitfront.io/r/user-7646844/ZTQB4rfx5SYN/CustomLLM/.

Q48

With which stakeholders are test results communicated?

☐ Compliance professionals

☐ Digital marketing professionals

☐ Model development professionals

☐ Production and maintenance professionals

☐ None

Fig. 5. 'Who...' questions

The overview of how the automated scoring process was implemented is shown in Fig. 7. Each respondent's score for each category is calculated along with the reflection score. The reflection score is a measure of how well an organization understands its own compliance with the AI Act. Figure 7 shows that this score is calculated by determining the ratio of "statement questions" and other questions (process). Statement questions rely primarily on respondents' perspectives, whereas other questions are more objective. The reflection score determines if an organization over- or underestimates itself.

Question	Answer	N. of questions	Scoring
Statement	1-6 scale	25	Strongly agree → 2 points Somewhat agree → 1 point Neutral → 0.5 point Disagree → 0 points
How often...	MC	24	For every model → 2 points For most models → 1 point For a few or never → 0 points
Do you...	Yes/No	14	Yes → 2 points No → 0 points
Who...	MC	9	0.5 point per stakeholder
To what extent...	0-10 scale	5	Not scored, but used to gain insight into the organization characteristics.
How many...	MC	4	If someone is responsible → 2 points If multiple people are partially responsible → 1 point If no one is responsible → 0 points
When...	MC	3	In every iteration of a project → 2 points Once at a relevant time → 1 points Sometimes or at the wrong time → 0 points
What type...	MC	2	Differs per question. Generally if a suitable method is present 2 points, if they are working on it 1 point otherwise 0 points.
Overig	MC	3	Differs per question.

Fig. 6. Points given per question type

For example, if an organization strongly agrees that they communicate accepted risks of the system with the user, but also states that they never measure a model's risk, there appears to be an overestimation by the respondent. Conversely, if an organization scores low on statement questions but high on other questions, it may be underestimating its compliance. The reflection score is added to show the validity of the responses and to help organizations better understand their own compliance.

4   Results

4.1   Identifying Focus Areas

The average percentage score for each category of the questionnaire is shown in Fig. 8. The overall average compliance score for all respondents and categories is 57%. The average reflection score is 1.0, suggesting that organizations demonstrate good self-awareness. Figure 8 reveals variations in compliance scores across different categories. The questionnaire results show that many organizations lack procedures for technical documentation and do not have someone trained to determine compliance requirements. Regarding data and model internals, organizations struggle with training employees on data and model bias. User communication presents challenges in determining metrics for measuring model risks on rights and discrimination. Risk management systems are found to be lacking in some organizations. Model monitoring shows a mixed trend, with some organizations adequately updating models when needed and having protocols in place to determine if data is outdated.

4.2   Organization Characteristics' Influence

Organizations with 1–50 employees scored lower than organizations with 51+ employees. As for the industry, the dataset is too small to draw any conclusions. There is a large variance in compliance scores for the IT sector, from 26% to 67%. One organization felt that their ISO certification helped them to comply with the AIA. Organizations with more AI experience in years did not score better compared to organizations relatively new to AI.
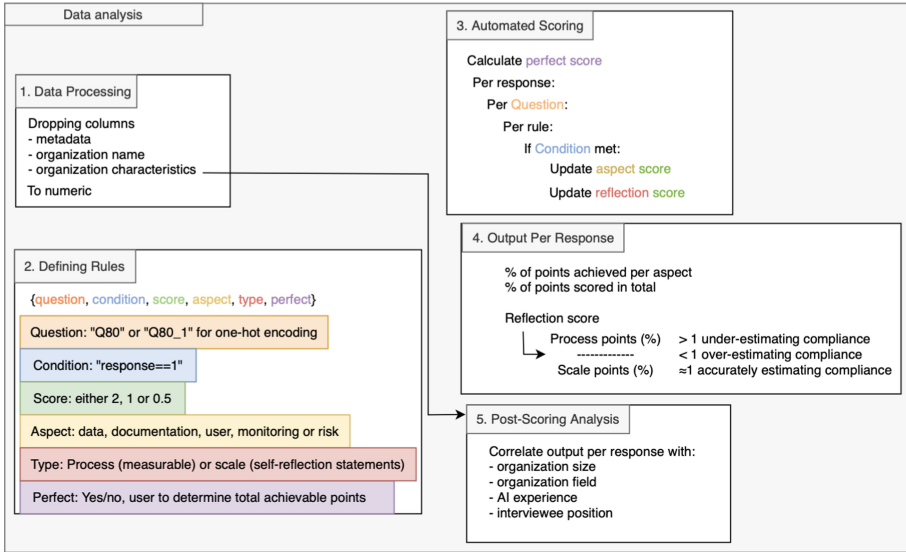
Fig. 7. Scoring system of questionnaire data

4.3   Prevalent Questions

Besides the challenges mentioned in Sect. 4.1, this paper also contributes to better understanding the challenges of organizations by identifying common questions among respondents. The interactive interviews identified several prevalent questions among respondents, both on the content of the AIA as on the application within their organization. These questions are either directly asked by the organization, or extracted from the answers on the questionnaire. For instance, when asked what data risks organizations have dealt with in the past years, almost all organizations gave an answer relating to GDPR compliance. In reality, there are many other risks besides GDPR compliance that the AIA is concerned with, so the question would be, 'What other risks besides data privacy should my organization be concerned with?'.

The identified prevalent questions among most organizations are as follows: Questions on the content of the AIA:

1. Should technical documentation also be written for non-technical people?
2. Does the AIA stipulate that we need someone to monitor the AI models full-time?
3. Does the AIA require me to work with encrypted data only?
4. How should we deal with missing data according to the AIA?
5. What other data risks besides data privacy should my organization be concerned with?
6. What does the AIA mean by high-risk AI?
7. Does the AIA require an external audit?

8. Which documents should be included in the compliance documentation?
9. Does the AIA mention metrics that should be used to determine a model's risks for rights and discrimination?
10. What does the AIA mean by 'human oversight'?

Questions on the application of the AIA within their organization:
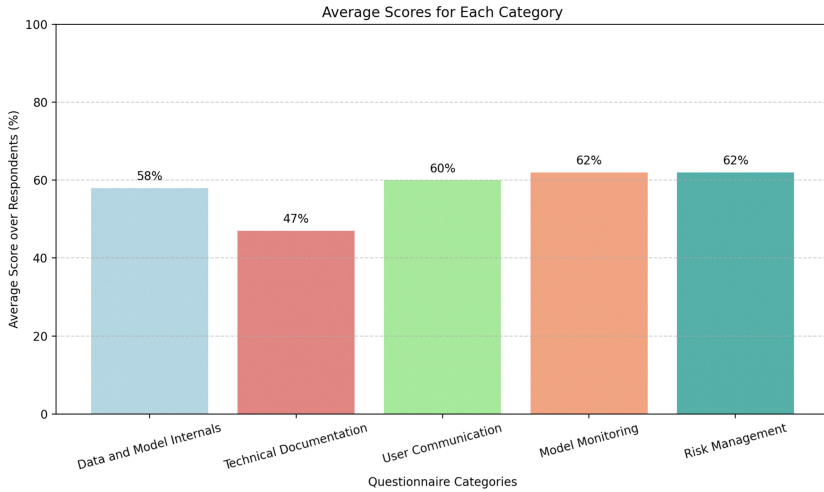
Fig. 8. Percentage of points per questionnaire category

1. To which extent does my ISO certification help towards AIA compliance?
2. Does GDPR training also include data bias and model bias training?
3. What are the biggest risks to AIA compliance when data is gathered in-house?
4. Our organization uses data from customers; what are some of the biggest risks when aiming for AIA compliance?
5. We only use ChatGPT and other out-of-the-box AI models; should we still be concerned with the AIA?
6. What can we do to improve AIA compliance concerning our technical documentation?
7. We currently don't communicate anything about our models with our users; how can we better communicate information with the users for AIA compliance?
8. Our organization is very small, and no one is specialized in compliance; where do we even begin to achieve AIA compliance?
9. We currently have no idea if we communicate with our stakeholders according to the AIA; how should we assess this to make improvements?
10. The AIA stipulates that accuracy should be according to the state of the art. This seems very vague; how should I go about achieving state-of-the-art accuracy?

These questions can be useful for future research to understand the needs of organizations.[3]

5   Conclusion

This paper examines in which areas organizations seem to be struggling with regards to current and future compliance with the AIA. A conceptual framework has been constructed based on a review of the act. A questionnaire is formulated based on the framework. Fifteen organizations answered the entire questionnaire.

A compliance score is calculated using a rule-based system that awards points for answers following the contents of the AIA. Organizations achieve an average compliance score of 57% compared to the 'perfect' score. This score indicates there is room for improvement towards AIA readiness. Organizations are best prepared on model monitoring and risk management but score the lowest, with 47%, on technical documentation.

The duration of AI usage by organizations does not result in a higher compliance score. The same goes for IT organizations compared to non-IT organizations. Overall, this paper contributes to the growing body of knowledge on the impl042weementation of the AIA. The paper is the first to identify focus areas for different categories of the AIA to help organizations better prepare. Organizations will need help dealing with the questions and challenges they are facing.

6   Future Research

Several approaches for future research are identified. First, the predictive power of the questionnaire should be tested to see if the questionnaire can predict if an organization will pass the self-assessment. More qualitative data should be gathered by observing the AIA compliance processes. This data can then be used to refine the questionnaire for different organizations' sizes and sectors.

---

[3] Future research from one of the authors of this paper has focused on examining how organizations can be supported with these questions: https://www.babelfish.nl/blog/unraveling-aia-with-llm.

References

1. The Act: The Artificial Intelligence Act. https://artificialintelligenceact.eu/the-act/. Version 21 Apr 2021
2. European Commission: Proposal for a regulation laying down harmonized rules on artificial intelligence (Artificial Intelligence Act) and amending certain EU legislative acts (2021). https://eur-lex.europa.eu/legal-content/TXT/?uri=CELEX%3A52021PC0206
3. Usman, M., Felderer, M., Unterkalmsteiner, M., Klotins, E., Mendez, D., Alégroth, E.: Compliance requirements in large-scale software development: an industrial case study. In: Morisio, M., Torchiano, M., Jedlitschka, A. (eds.) PROFES 2020. LNCS, vol. 12562, pp. 385–401. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64148-1_24
4. da Conceca Freitas, M., Mira da Silva, M.: GDPR compliance in SMEs: there is much to be done. J. Inf. Syst. Eng. Manag. (2018). https://doi.org/10.20897/jisem/3941
5. Chua, H.N., Herbland, A., Wong, S.F., Chang, Y.: Compliance to personal data protection principles: a study of how organizations frame privacy policy notices. Telematics Inf. (2017). https://doi.org/10.1016/j.tele.2017.01.008
6. Veale, Zuiderveen Borgesius, F.: Demystifying the draft EU artificial intelligence act. CRi Comput. Law Rev. Int. (2021). https://arxiv.org/pdf/2107.03721.pdf
7. Ebers, M., Hoch, V.R.S., Rosenkranz, F., Ruschemeier, H., Steinrtter, B.: The European commissions proposal for an artificial intelligence act critical assessment by members of the robotics and AI law society (RAILS). J. **4**(4), 589–603 (2021). https://doi.org/10.3390/j4040043