Legal Knowledge and Information Systems
G. Sileno et al. (Eds.)

© 2023 The Authors.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/FAIA230955

Automating Business Process Compliance for the EU AI Act

Claudio NOVELLI [a] Guido GOVERNATORI [b] Antonino ROTOLO [a]
[a] *Centre for Digital Ethics - DSG, and Alma AI, University of Bologna, Italy*
[b] *Brisbane, Australia*

Abstract. The EU AI Act is the first step toward a comprehensive legal framework for AI. It introduces provisions for AI systems based on their risk levels in relation to fundamental rights. Providers of AI systems must conduct Conformity Assessments before market placement. Recent amendments added Fundamental Rights Impact Assessments for high-risk AI system users, focusing on compliance with EU and national laws, fundamental rights, and potential impacts on EU values. The paper suggests that automating business process compliance can help standardize these assessments and outlines some methodological guidelines.

Keywords. Risk Assessment, AI Act, Business Process Compliance

1. Introduction

The EU AI Act (AIA) is the first initiative towards a comprehensive legal framework on AI. It is a very ambitious legislative project, which however leaves open some issues and implementation problems. As is well-known, AIA categorises AI systems (AIs) into four risk categories—unacceptable, high, limited, and minimal—assigning corresponding regulatory and procedural burdens to their providers and deployers.

The most challenging category of AIs is high risk [9], because it covers domains such as biometric systems, critical infrastructure, education, administration of justice and democratic processes, where serious risks are possible, but where also potential advantages for individuals and for the public interest are possible as well [8]. Attention has been so far devoted to conformity assessment procedures (CA), which consist in a risk-assessment analysis that providers of AIs must perform before they place such systems on the market. From the business perspective, [7] suggests to implement at AI organisations platforms ensuring computational accountability, which also specify organisational measures and computational methods to standardise CA procedures.

While the implementation of CA is still an open question, another issue needs to be addressed almost from the scratch. In fact, recent amendments by EU Parliament have expanded for deployers of high-risk systems transparency obligations, which complement the Commission's previous focus on providers. Article 29 introduces new responsibilities for deployers, emphasizing the protection of individuals' fundamental rights in AI decision-making. They must notify individuals when these decision-making systems are applied, clarify their intended purpose, and specify the types of decisions involved. Article 68c grants individuals the right to an explanation for AI-generated decisions, enhancing transparency in decision-making involving AI. Article 29a introduces a *Fun-*

damental Rights Impact Assessment (FRIA) before deploying high-risk AIs, akin to a Data Protection Impact Assessment (DPIA) for the GDPR.

How can the FRIA be concretely developed? To the best of our knowledge, no technical solution has been so far elaborated in the literature, except a checklist developed by the Dutch Government[1]. This paper aims at offering some ideas to fill the gap and some guidelines for FRIA in the perspective of standardising and automating it.

The layout of the paper is as follows. Section 2 accounts for the FRIA, recalls a model for risk assessment proposed elsewhere, and identifies some guidelines for the automation of FRIA. Section 3 illustrates how Business Process Compliance methods can be used for the FRIA.

2. Four Guidelines for the Fundamental Rights Impact Assessment

The amended AIA mandates that *all users* of high-risk AI systems conduct a FRIA. Users must develop a detailed plan to mitigate negative fundamental rights impacts or inform the AI provider and national authorities promptly. The idea behind the FRIA sees the AI risk as a potential harm resulting from AI violating legal values and requires assessing the potential (a) impact of AIs on the compliance with obligations applying to them, (b) detriment of fundamental rights determined by the deployment of AIs.

In [8] we argued that we cannot treat values as technical standards, leading to predetermined outcomes in the balancing test of values and interests, and without flexibility for risk management adjustments based on changing circumstances. Otherwise, we would result in an inaccurate evaluation of AI risk. We rather have to shift from a purely scope-oriented categorisation of AI risks to an analysis based on risk scenarios involving interactions among multiple risk factors.

To do so, we have proposed to adapt in the context of AIs the risk assessment model arising from the Intergovernmental Panel on Climate Change (IPCC) and related literature. This integrated model enables the estimation of AI risk magnitude by considering the interaction between (a) risk determinants, (b) individual drivers of determinants, and (c) multiple risk types [8]. Risks is the consequence of hazard, exposure, and vulnerability. Hazard refers to potential sources of harm. Exposure refers to what might be affected by the hazard source. Vulnerability refers to attributes or circumstances that make exposed elements susceptible to harm.

The interactions among risk factors determine the two input variables of the overall risk magnitude:

- the likelihood of the event depend on the interaction between hazard drivers and response drivers (e.g., preventive measures);
- the severity of the detriment can be higher or lower depending on the hazard sources, exposed asset, and vulnerability profiles.

The second step is to evaluate the suitability of the resulting risk assessment in relation to the asset exposed to the use of an AIs, by means of a *process of balancing of fundamental rights and values* relative to a given deployment context. Intuitively, this is needed because the potential detriment of some rights can be balanced by the promotion of other values.

[1] https://www.government.nl/documents/reports/2021/07/31/impact-assessment-fundamental-rights-and-algorithms

Gaps in liability that arise when AIs are used have an economic impact. In fact, legal compliance and liability rules would minimise costs of harm related to AIs. However, compliance carries a cost as well, and without, e.g. proportionality judgement in the FRIA—as we argued in [8]—the AIA may become unsustainable for AI deployers (and providers). This could jeopardise the EU's AI strategy, hinder innovation, and miss the potential benefits AI aims to protect. Therefore, the AIA requires a clear risk assessment model whose implementation should follow some general design guidelines:

Guideline 1 (Standardisation). *Define business processes for FRIA that are integrated with other standardised business policies. For example, follow the design principles adopted for managing the DPIA from GDPR.*

If FRIA can be standardised, the next step is its automation, which would certainly reduce costs and, above all, make the FRIA itself more robust, scalable, and reliable.

Guideline 2 (Automation of Regulatory Compliance of FRIA). *Formalise FRIA compliance processes within existing computation methodologies that support compliance checking in the field of Business Process Management.*

For the automation of FRIA, one could argue that the legal risk assessment can profit from the data stored in the AI public registry. However, machine learning techniques are not the best option, because they are indisputably AI systems, and they lack full transparency and explainability: if used, the consequence could be that the AI system in support of FRIA is in turn a high-risk AI system needing another FRIA.

Guideline 3 (Transparent Automation of FRIA). *Adopt transparent and explainable computational methods for implementing the FRIA.*

Finally, the automation for FRIA should be be comprehensive:

Guideline 4 (Full automation). *Adopt computational methods for implementing the FRIA for a system X that cover (a) the identification of obligations relevant for X and the automatic compliance monitoring; (b) procedures for balancing the involved legal values.*

3. Business Process Compliance for the FRIA

Business Process Compliance [4] is a methodology that proved successful in supporting compliance checking in the field of Business Process Management. Here we outline the methodology of [6] and advance how to use it for the AIA.

A business process model describes the activities and the order in which they are typically executed by an organisation to achieve a business goal. A *trace* of a business process model is a sequence of tasks in the process that adheres to the order and constraints specified by the model. Consider the process in Figure 1, in standard BPMN notation, where we have a task A followed by an XOR split; this means that after A we have two ways to proceed. We have the choice to perform B or C. In the XOR-split in one of the branches we have task B followed by the AND-split of a branch with task D, and a branch consisting of only task E; accordingly, we have to execute the tasks in the branches following the order in the branches. The second branch of the XOR-split has only one task: C. The traces of the process are $\langle A, C\rangle$, $\langle A, B, D, E\rangle$ and $\langle A, B, E, D\rangle$. Given a process P, $\mathcal{T}_P = \{t_1, t_2, \ldots\}$ denotes the set of traces of P.

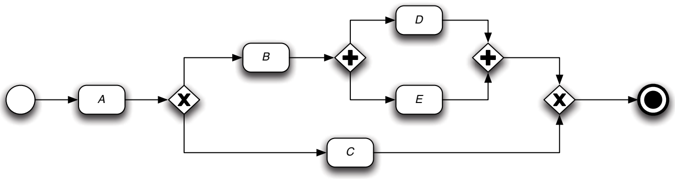

Figure 1. Example of a business process model in standard BPMN notation

Section 2 identified two compliance layers for the FRIA: (a) ensuring and checking the regulatory compliance with obligations applying to AIs; (b) assessing the potential detriment of fundamental rights determined by the deployment of AIs.

3.1. FRIA: Compliance with Obligations

In the best situation (e.g., process aware information system), business process models of AIs are available, and then our methodology can be applied directly. In the most common situation business process models are not available, and one might argue that in some case they do not even exists (for example for machine learning based approaches). Fortunately, the Act itself can help here, since it mandates that (high risk) AI approaches have to automatically generate logs of the activities they perform. Here, we can use another methodology originated from business process management: process mining [10]. The idea behind process mining is to use process logs to extract the information about what activities, data, resources are used by a computer system, and the ordering relation on such information. The aim is to use the extracted information to provide a(n approximate) process model corresponding to the log (and then the activities performed by the system).

Compliance is not only about the tasks, but it is concerned also on their effects (i.e., how the activities in the tasks change the environment in which they operate), and the artefacts produced by the tasks (for example, the data resulting from executing a task or modified by the task) [5]. To capture this aspect process models are enriched with semantic annotations for tasks. An annotation is a set of formulas giving a (partial) description of the environment in which a process operates.

To model the process' states we use the function

$$State\colon \mathcal{T}_P \times \mathbb{N} \mapsto 2^{\mathcal{L}},$$

where $\mathcal{L}$ is the set of formulas of the language used to model the annotations. Let us illustrate with an example the meaning of the function *State*. Suppose we have the trace $t = \langle A, B, D, E\rangle$, and that $State(t,3) = \{p,q,r\}$. This means that $\{p,q,r\}$ is the state resulting after executing D in the trace t (D is the third task in t). A trace uniquely determines the sequence of states obtained by executing the trace. Thus, we use a trace to refer to a sequence of tasks, and the corresponding sequence of states.

Norms produce legal effects, such as *obligations*, and are constraints that limit the space of action of processes. Compliance means to identify whether a process violates or not a set of obligations. Thus, the first step is to determine whether and when an obligation is in force. Hence, an important aspect of the study of obligations is to understand the lifespan of an obligation and its implications on the activities carried out in a process. [6] provides a comprehensive classification of the obligations in terms of their life-cycle from a compliance point-of view. Defeasible Deontic Logic (DDL) [3], for example, supports all deontic notions in [6] and has mechanisms to terminate and remove obligations [2]. We introduce the function

$$Force\colon \mathcal{T}_P \times \mathbb{N} \mapsto 2^{\mathcal{L}}$$

that, given a process P associates to each task in a trace a set of literals, where these literals represent the obligations in force for that combination of task and trace.

The interaction between *State* and *Force* determines if the state of a process after the execution of a task in a trace results in a breach of the norms governing the process.

The set of traces of a given business process describes the behaviour of the process insofar as it provides a description of all possible ways in which the process can be correctly executed. Accordingly, for the purpose of defining what it means for a process to be compliant, we will consider a process as the set of its traces.

Intuitively a process is compliant with a set of normative constraints, which we call *normative system*. Two notions of compliance can be defined:

Definition 1. *Let $\mathcal{N}$ be a normative system, and P a process.*

1. P fully complies *with $\mathcal{N}$ iff every trace $t \in \mathcal{T}_P$ complies with $\mathcal{N}$.*
2. P partially complies *with $\mathcal{N}$ iff there is a trace $t \in \mathcal{T}_P$ that complies with $\mathcal{N}$.*

Definition 2. *A trace t complies with a normative system $\mathcal{N} = \{n_1, n_2, \dots\}$ iff all norms in $\mathcal{N}$ have not been violated.*

The problem of determining whether a business process is compliant amounts to populate *State* and *Force* function. Since the number of states in a process grows exponentially, norms must be formalised in efficient formalisms such as DDL. In addition, a domain use can annotate the tasks of the process with the formulas (from the vocabulary of the norms) that corresponds to the effects of the tasks. Once the norms are formalised and the process is annotated we can use the procedure of [2] to check compliance.

3.2. FRIA: Impact on Fundamental Rights

The second layer of FRIA concerns assessing the potential detriment of fundamental rights determined by the deployment of AIs. We argued that a proportionality judgement must be modelled in order to check, e.g., if the relative negative impact of an AI with respect to legal value v_1 is balanced by the promotion of another value v_2. We outline a methodology that can be integrated in the one described in the previous section.

Let $\mathcal{V} = \{v_1, v_2, \dots\}$ be a set of legal values for the AIA. Following [1], values can be ordered using the operator $\otimes$. The interpretation of an expression $v_1 \otimes v_2$ is that v_1 is the most preferred value, but, if v_1 is demoted then v_2 is preferred. Such an ordering can be established by the developers of the FRIA implementation or are extracted from the AI registry or from legal practice.

As appropriate with legal balancing, orderings on values are relativised to contextual conditions $a_1, \dots, a_n$, which in DDL can be expressed through rules like:

$$a_1, \dots, a_n \Rightarrow v_1 \otimes \cdots \otimes v_n \tag{1}$$

Any context C is meant to capture the specific deployment conditions of an AIs, i.e., the environment in which a process operates and the effects of deployment.

We can establish if any C *promotes* or *demotes* some values:

$$r_1 : a_1, \dots, a_n \Rightarrow \mathsf{Promotes}(v_j) \qquad r_2 : b_1, \dots, b_n \Rightarrow \mathsf{Demotes}(v_k)$$

where $a_1, \ldots, a_n, b_1, \ldots, b_n \subseteq C$. The basic reasoning mechanism works as follows: given a context $C = \{a_1, \ldots, a_n\}$

- determine the set $\mathbf{P}$ of value preferences $\{\bigotimes_{i=1}^{n} v_i, \bigotimes_{k=1}^{m} v_k, \ldots\}$ holding in the context C using rules such as (1);
- establish through rules like r_1 and r_2 which values are promoted and which ones are demoted;
- define *the degree of value compliance* $\mathbf{Degree}(C)$ *of* C as

$$\mathbf{Degree}(C) = \left[\sum_{x:\mathsf{Promotes}(v_x)} \left(\frac{\sum |\{v_x | \bigotimes_{x=1}^{n} v_x \in \mathbf{P}\}|}{x}\right)\right] - \\ - \left[\sum_{y:\mathsf{Demotes}(v_y)} \left(\frac{\sum |\{v_y | \bigotimes_{y=1}^{m} v_y \in \mathbf{P}\}|}{y}\right)\right]$$

A context C is *optimal* iff $\mathbf{Degree}(C)$ is maximal or iff it is greater than a certain acceptability threshold. The degree of value compliance is an evaluation of the suitability of the resulting risk category in relation to the asset exposed to the use of an AIs.

4. Summary

Recent amendments of AIA have introduced Fundamental Rights Impact Assessments, which are addressed to all users of high-risk AI systems. This level of assessment is risk-based and requires checking the compliance with EU and national legislation and with fundamental rights law, as well as considering the potential negative impact on EU values and rights.

We offered in this paper some ideas on whether techniques for the automation of business process compliance can help users in standardising such types of assessment. We argued that this is possible, and some guidelines and methods are described. Such methods extend existing algorithms for ensuring or checking business process compliance.

References

[1] Erica Calardo, Guido Governatori, and Antonino Rotolo. Sequence semantics for modelling reason-based preferences. *Fundam. Informaticae*, 158(1-3):217–238, 2018.

[2] Guido Governatori and Antonino Rotolo. A conceptually rich model of business process compliance. In *APCCM 2010*, CRPIT 110, pages 3–12. Australian Computer Society, 2010.

[3] Guido Governatori, Antonino Rotolo, and Giovanni Sartor. Logic and the law: Philosophical foundations, deontics, and defeasible reasoning. In *Handbook of Deontic Logic and Normative Reasoning*, volume 2, pages 655–760. College Publications, London, 2021.

[4] Guido Governatori and Shazia Sadiq. The journey to business process compliance. In *Handbook of Research on BPM*, chapter 20, pages 426–454. IGI Global, 2009.

[5] Mustafa Hashmi, Guido Governatori, and Moe Thandar Wynn. Business process data compliance. In *RuleML 2012*, LNCS 7438, pages 32–46. Springer, 2012.

[6] Mustafa Hashmi, Guido Governatori, and Moe Thandar Wynn. Normative requirements for regulatory compliance: An abstract formal framework. *Information Systems Frontiers*, 18:429–455, 2016.

[7] Joris Hulstijn. Computational accountability. In *ICAIL 2023*, ICAIL '23, page 121–130. ACM, 2023.

[8] Claudio Novelli, Federico Casolari, Antonino Rotolo, Mariarosaria Taddeo, and Luciano Floridi. Taking AI Risks Seriously: a New Assessment Model for the AI Act. *AI & Society*, 38, 2023.

[9] Jonas Schuett. Risk management in the artificial intelligence act. *European Journal of Risk Regulation*, page 1–19, 2023.

[10] Will van der Aalst. *Process Mining: Data Science in Action.* Springer, 2016.