Challenges in Digital Compliance: Risk Assessment and Fundamental Rights under the GDPR and the EU AI Act

Olga Kokoulina[1]

[1]Center of Private Governance (CEPRI), Faculty of Law, University of Copenhagen, Karen Blixens Plads 16, 2300 Copenhagen, Denmark

Abstract

Recent EU legislative developments, particularly the General Data Protection Regulation (GDPR) and the AI Act have significantly propelled the integration of risk management within the digital compliance realm. Risk, being a multifaceted concept, permeates a vast array of policy domains, extending from traditional areas such as environmental and safety regulations to domains as varied as housing and child protection. It underpins policy formulation and enforcement prospects and is often deeply enmeshed with considerations of "harm" and "liability". Taking a risk assessment as a case study for digital compliance, the discussion probes into the challenges and constraints of automating such processes. In particular, it investigates the feasibility of translating references to "risks to fundamental rights and freedoms of individuals" as outlined in the compliance requirements of the GDPR and AI Act, into digital frameworks.

Keywords

risk, digital compliance, GDPR, AI Act

1. Introduction

Complying with the requirements of EU data-related legislation clearly underscores the dire necessity of an interdisciplinary approach; such engagement is not only desirable but also largely anticipated. For instance, the GDPR imposes multiple obligations on data controllers to adopt appropriate technical measures that fulfill their responsibilities, uphold the rights of data subjects, and ensure the overall implementation of the Regulation. More often than not, these requirements serve as general guidelines for meeting specific data subjects' rights and data controllers' obligations. However, in certain cases, such as within the context of the Data Protection Impact Assessment (DPIA), these requirements may come as directly tied to clearly specified compliance mechanisms (Article 35). In addition to emphasizing the adoption of appropriate technical measures, the GDPR also establishes overarching principles of data protection by design and by default. The core objective of these concepts is to ensure that data protection considerations are integrated at the system design level. While data protection by design involves embedding data protection principles into the overall design of the data processing framework, data protection by default focuses primarily on adhering to data minimization and proportionality principles.

Thus, the GDPR includes several examples where potential collaboration with computer science naturally emerges as an effective way to implement the regulatory requirements. A similar observation applies to the recently adopted EU AI Act, where the need for interdisciplinary collaboration, particularly with computer science, is also evident. For instance, as the Act categorizes AI systems based on their risk level, it further adjusts the obligations framework to align with the specific level of anticipated risk. One of the key elements of this framework is a requirement to conduct a specific – fundamental rights oriented - assessment in particular cases where employment of high-risk AI systems is expected to affect the rights of individuals or groups of individuals (e.g. Recital 96, Article 27). The purpose of this Fundamental Rights Impact Assessment (FRIA) is to proactively engage the deployer of the high-risk AI systems in the identification and mitigation of the potential risks in a continuous manner: before the product in question is placed on the market, as well as at a later stage, shall any relevant factors

PLC - Processes, Laws and Compliance workshop, in conjunction with ICPM 2024, October 14, 2024, Lyngby, Denmark

olga.kokoulina@jur.ku.dk (O. Kokoulina)

influencing such assessment change (Recital 96). Standardization and automation of the process of such assessment can potentially bring about significant efficiency gains, and, at least theoretically, pose a great promise for a synergy between legal and computer science communities.

While both the GDPR and the AI Act offer significant opportunities for collaboration, they also underscore the undeniable challenges of identifying and interpreting the proper knowledge bases for automation. The challenges particularly involve the long-identified difficulty of appropriate knowledge scoping, translating ambiguity inherent in open-textured and context-dependent concepts, as well as the ongoing need to ensure that systems are continuously revised and remain flexible to keep pace with evolving societal views and interpretations. Against this backdrop, the discussion aims to advance scholarly engagement by merging a theoretical exploration of risk assessment's scope within normative models with a practical examination of the challenges associated with automating these processes.

2. "Risk" under the GDPR and AI Act as a case study

Firstly, it is important to establish a common understanding of the "risk" as a concept. In the legal domain, "risk" has long been employed as a governance tool, even though its application has varied across different regulated sectors and geographic regions[1]. Having said that, one cannot help but notice an increasing recourse to "risk" as an often-preferred means for setting EU regulatory and enforcement priorities, particularly in the data-driven economy and in the context of emerging technologies. For example, although the concept of a 'risk-based approach' was already embedded in the earlier European data protection law—Directive 95/46— particularly in its provisions on the processing of special categories of data (Article 8), security requirements (Article 17), notification procedure (Article 18) and prior-checking obligations for data protection authorities (Article 20), its successor, the GDPR, incorporates and emphasizes the notion of risk much further and more extensively.

The GDPR, much like Directive 95/46, uses the concept of 'risk' to classify certain types of data as information requiring a heightened level of protection. In this paradigm, "risk" functions as an inherent characteristic, a built-in aspect of the notion itself. Certain types of data are automatically deemed to carry elevated "risk" simply due to their nature and sensitivity. It follows, the concept of risk is not something to be evaluated on a case-by-case basis in the case of sensitive data. Rather, it is assumed as an intrinsic quality of processing of the specific data categories: such processing is prohibited in principle, unless specific circumstances and exceptions.

In other GDPR contexts, however, the concept of "risk" appears to align more closely with the idea of a "scalable and proportionate approach to compliance"[2], particularly from the perspective of those bearing the compliance obligations. The core premise is that data subjects should receive a consistently high and effective level of protection, regardless of whether the processing is deemed "low risk." However, a data controller involved in lower-risk processing may face fewer compliance obligations compared to the one handling higher-risk activities"[2]. This approach employs risk as a yardstick for tailoring safeguards and compliance measures on a case-by-case basis. It is especially prominent in GDPR provisions concerning the responsibilities of the controller (Article 24), requirements of data protection by design and by default (Article 25), security and related data breach obligations (Articles 32 and 34), and data protection impact assessments (Article 35). Despite referencing risk as a benchmark across these areas, the GDPR, however, lacks a statutory definition of 'risk.' Instead, it explicitly mentions it in two interconnected areas: first, as a cumulative criterion for determining the appropriate measures to be implemented, and second, as a requirement for undertaking a data protection impact assessment.

In the first scenario, "risk" is referenced alongside factors like the nature, scope, context, type, and purposes of processing, as well as considerations of the 'state of the art' and implementation costs. The wording of the respective articles indicates that the risk serves as one of the key determinants in assessing the appropriateness of technical and organizational measures towards demonstrating accountability, fulfilling data protection by design and by default, and meeting security obligations. A literal reading of the relevant provisions (Articles 24, 25, 32, 34) cautions against prioritizing any single factor: rather,

the corresponding sentences layer several considerations to be evaluated together. Such interpretation is also supported by references to the likelihood and severity of risks to the rights and freedoms of the data subjects, which are determined by evaluating the nature, scope, context and purposes of the processing (Rec.76). That said, it is also important to acknowledge that each of these provisions focuses on a specific dimension of risk, as reflected in the purpose of the measures they require. For instance, Article 24 emphasizes the implementation of suitable technical and organizational measures aimed at demonstrating overall compliance with GDPR requirements. Article 25 places significant emphasis on integrating data protection measures to proactively address compliance concerns from the very beginning, during the design and development phase of systems. Article 32, on the other hand, centres on the requirement to establish security measures that safeguard the confidentiality, integrity, availability, and resilience of data [3].

The second context pertains to data protection impact assessments (DPIA). Unless the processing activities in question fall under specific exceptions (Art.35(5) and Art.35(10)), such assessments are required if data controllers determine that the processing is likely to result in a high risk to the rights and freedoms of individuals. The outcome of this assessment, in its turn, should be again instrumental for determining the appropriate measures to be taken in order to demonstrate compliance (Rec. 84) and mitigate identified risks (Rec.90). Unlike provisions in the first context, where risk is treated as one of several principal and interconnected factors to consider, the DPIA context specifically focuses on situations involving "high risks." Much like the GDPR's distinction between "ordinary" data processing in Article 6 and the processing of the special category data in Article 9, Article 35 acknowledges that high-risk processing requires more than just the baseline measures; it necessitates an elevated level of scrutiny and precaution for high-risk processing activities. Article 35(3) of the GDPR provides for certain high-risk scenarios that trigger the need for a DPIA, including processing activities such as evaluation of individuals' personal characteristics based on automated processing, large-scale processing of sensitive data or of personal data relating to criminal convictions, and systematic monitoring of publicly accessible areas. As suggested by the expression 'in particular' in the introductory sentence, this list of high-risk processing activities is intended to be non-exhaustive. Furthermore, Article 35 (1) specifically highlights the use of new technology as one of the key criteria that may necessitate heightened consideration for conducting a DPIA. Thus, according to Article 35(1) and recitals 89 and 91, the employment of new technologies for data processing can trigger a DPIA requirement due to the potential for novel data collection and usage practices that may pose significant risks to individuals' rights and freedoms.

From a digital compliance perspective, this second context is especially compelling for several reasons. To begin with, it is based on a relatively defined set of circumstances illustrated by examples within the GDPR itself, as well as in the guidelines provided by the Article 29 Working group/European Data Protection Board and national data protection authorities (Art.35(4). Some supervisory authorities have also published lists of specific types of processing operations for which DPIAs are not required. Therefore, considering the potential scope of the knowledge base for automation, there is relatively greater clarity and predictability regarding which types of data activities necessitate a more comprehensive impact assessment. While the focus on the impact on human rights in Articles 24, 25, and 32 is supplemented by additional contextual requirements for risk assessment and tends to be tailored to the specific needs and circumstances of each organization, the DPIA requirements in Article 35 appear, at least at first glance, to be more broadly applicable across a variety of organisations.

Furthermore, The DPIA offers strong potential for automation, as it can serve as a reusable blueprint for similar and subsequent data processing assessment, as follows from Article 35 (1). Thus, building a process to run the assessment prior to data processing allows organizations to conduct a single assessment for multiple similar operations, enhancing efficiency and consistency in compliance efforts.

Finally, there is a notable similarity between the data processing activities that require a DPIA under the GDPR and the high-risk AI applications that demand fundamental rights assessments under the AI Act. For instance, both frameworks identify biometric data processing, such as facial recognition, public surveillance, and profiling as high-risk due to their implications for privacy and freedom. Similarly, the use of AI in critical areas like employment, education, access to essential services, and law enforcement

is categorized as high-risk under the AI Act, aligning with GDPR's requirement for DPIAs in data processing that influences employment decisions, educational access, or eligibility for social benefits. The overlaps between GDPR's DPIA requirements and the AI Act's high-risk assessments present valuable opportunities to leverage shared criteria and integrate technological tools into compliance processes, potentially making both assessments more efficient. Similar to the DPIA requirements under the GDPR (Article 35(7), and recitals 84 and 90), the AI Act also provisionally outlines minimum requirements for the content of the risk assessment. Additionally, it stipulates that national AI offices are expected to develop templates and automated tools to streamline and support compliance efforts. Thus, the experience of conducting and potentially automating DPIA assessments under the GDPR is valuable in ways that could be extendable and complimentary to compliance efforts under the AI Act.

However, upon closer examination, the automation of compliance assessments in both cases is fraught with potentially significant challenges.

3. Automating Risk Assessment

As noted, the GDPR does not explicitly define the concept of "risk." Beyond the pointed-out contextual examples and instances of data processing activities that are considered "high risk," the Regulation offers limited direct guidance on how "risk" – as a notion – should be conceptualized. From an automation perspective, in principle, this omission does not come across as an insurmountable obstacle: the examples of data processing activities that trigger the obligation to conduct a DPIA offer valuable criteria for building a roadmap to determine whether certain activities fall under this requirement. These examples provide a framework that can guide automated systems in identifying high-risk activities by mapping them against established risk indicators. By embedding these criteria, an automated tool could efficiently check whether specific processing operations align with those listed as high-risk, helping organizations consistently and methodologically assess their need for a DPIA.

However, this approach should also align with the additional guidance provided in Recital 75, which has not yet been considered. The recital in question seems to suggest that a foundational understanding of "risk" could be derived from analyzing the potential harm that a particular data processing activity could cause. Thus, risks to individuals may arise from data processing that could lead to physical, material, or non-material damage. Reflecting on the potential damages that processing activities might entail—such as discrimination, identity theft, fraud, loss of control over personal data, or other significant economic or social disadvantages—calls for a deeper approach than merely checking off a list of high-risk activities. Notably, despite the potentially far-reaching implications of this interpretation of the risk, the guidance from the Article 29 Working Party does not delve into the specific content of what this deeper approach should entail[4]. So, on one hand, there is a relatively clear area of guidance as to what type of data processing activities merits closer scrutiny and the respective consideration in the DPIA. On the other hand, the GDPR also suggests a more nuanced analytical approach that goes beyond merely identifying high-risk activities, encouraging organizations to assess potential "harm" as part of their risk evaluation.

A similar approach is evident in the AI Act: unlike the GDPR, the AI Act provides an explicit definition of "risk" as "the combination of the probability of an occurrence of harm and the severity of that harm" (Art. 3). This definition establishes a clear emphasis on "harm" as a fundamental benchmark for risk assessment. By explicitly framing risk in terms of both the likelihood and gravity of potential harm, however, the AI Act stops short of providing an autonomous definition of it, leaving the concept open to interpretation and requiring organizations to apply context-specific judgments.

In practice, this approach could be challenging to implement within the efforts to automate compliance. Assessing risk based on the probability and severity of harm—as required by the AI Act and implicitly suggested by the GDPR—necessitates a nuanced understanding of the context, individual rights, and potential impacts, which can be difficult to fully capture in automated systems. Automated compliance tools may struggle to accurately gauge subjective elements like harm severity or to interpret complex, context-dependent risks that require human judgment.

This broader perspective moves beyond the conventional focus within technical research on risk mining, assessment, and analysis. In this field, research often aligns with the business management concepts of risk, defined as the ”effect of uncertainty on objectives,”[5],[6] as well as a “combination of the probability of an [unwanted] event and its consequences” [7].Consequently, risk identification functions as a key analytical step aimed at identifying factors that may obstruct the achievement of established objectives.

In recent years, efforts have been made towards leveraging automation to lay the foundation for defining these factors. Respective initiatives aim to streamline tasks such as data gathering, pattern recognition, and even preliminary risk assessments, allowing analysts to identify and assess potential risks and their impacts more efficiently [8][9].For instance, applying web mining and information extraction techniques has shown to be a promising way for developing a hierarchical risk taxonomy. Such a taxonomy could distinguish between specific, tangible risks a business might encounter, like interest rate fluctuations, and broader, abstract risk categories that encompass multiple specific types. Consequently, the framework could serve as a valuable computational tool to support risk-based decision-making in the financial sector[10].Another example could be found in investigating the capacity of sentiment analysis to predict financial risk. Related research undertakings explore the relationships between textual data and financial risk, assessing how qualitative “soft information” (texts in the form of, for example, opinions and market commentaries) might influence future risk analysis for companies[11].The common thread between these examples is the effort to transform available, often unstructured information into valuable insights for risk prediction, with the key point being that the data needed for such undertakings is accessible.

In the context of human rights impact assessment, however, obtaining comparable information about potential risks can be challenging. As a starting point, the previously discussed list of high risk data activities and high-risk AI systems offers some initial contours to help shape the foundational knowledge base. However, when it comes to examples of DPIAs under the GDPR or FRIA under the AI Act, there is no legal requirement to make these assessments publicly available, leaving limited material for analysis or mining. Relatedly, the specific scope of human rights included in such assessments remains somewhat ambiguous, further complicating efforts to establish a comprehensive framework. The Article 29 Working Party’s statement on the role of a risk-based approach in data protection frameworks suggests that while the phrase “the rights and freedoms” of data subjects primarily pertains to rights of data protection and privacy, it may also extend to other fundamental rights, such as freedom of speech, freedom of thought, freedom of movement, the prohibition of discrimination, and rights to liberty, conscience, and religion. Similarly, the scope of relevant rights for assessment under the AI Act remains unclear until further clarifications are provided. To assist with compliance, the AI Office is expected to develop a template for a questionnaire, possibly with an automated tool, aimed at helping deployers fulfill their obligations under the Act (Art. 27(5)).

Furthermore, while financial risks are still challenging to define with precision, there are nonetheless some empirical measures (e.g. volatility) and widely understood consequences, such as loss of funds, that provide a basis for assessment. In contrast, identifying risks to human rights and measuring impacts on individuals in terms of harm is even more complex, as it tends to be highly contextual and subjective, further complicating efforts to establish clear metrics. In other words, setting a baseline for risk assessment and defining the ultimate objectives that need safeguarding could be an insurmountable impediment to automation.

Finally, defining the objectives and identifying factors that might hinder their achievement (or ”risks” within this framework) is, in many respects, a normative exercise. This process requires interpretive judgment and raises a fundamental question: where does this competence reside, and how might it be effectively distributed and shared. For instance, the DPIA Guidelines emphasize that impact assessments should be grounded in the perspective of data subjects rather than organizations, unlike assessments in fields such as information security, which often prioritize organizational interests. In this context, it is essential to ensure that values and interests are genuinely considered, rather than regarded as implicit assumptions. Relatedly, as technology and data processing practices evolve, so too may the associated risks develop and change. It is therefore critical to ensure that the technology-supported assessment

is reiterative, flexible, and capable of adaptation to account for ever-changing new applications of technology and acknowledgement of unforeseen risks. Essentially, this requirement is central to the regulatory and enforcement challenge: when properly implemented, it can enhance the accountability of the compliance mechanism. If not, it risks reducing compliance to a meaningless checklist exercise.

4. Concluding Remarks

The concept of human rights impact assessment is not new [12], [13], and significant resources have been dedicated to operationalizing legal requirements, both in terms of internal management structures and practical implementation frameworks. The work in this area draws from multiple sources, including frameworks developed by national human rights institutions [14], Data Protection Authorities [15],[16], academic research [17], [18], and industry-led initiatives. The emerged foundation supports the potential for synthesis and further development through two key venues: stakeholder engagement and expert-driven project leadership.

Both the GDPR and the AI Act present opportunities for stakeholder engagement, facilitating input from a broad array of experts and affected parties to strengthen the legitimacy and thoroughness of such assessments (e.g. Recital 96 AI Act; Art. 35(9) GDPR).

Legal engagement in the development of such tools is crucial to ensure that regulatory expectations are incorporated as thoroughly as possible and that the assessment maintains a focus on potentially impacted individuals, rather than defaulting to an organization-centered perspective.

Equipped with this input, advancements in automation could elevate human rights assessments beyond basic questionnaires or matrix templates, leading to more sophisticated and dynamic tools capable of deeper analysis and predictive capabilities. This progression would be necessarily supported and scrutinized through the acknowledgement of implicit assumptions and the articulation of inter-disciplinary concerns. Nonetheless, this approach appears to offer an effective method for integrating diverse perspectives and expertise, ultimately creating a comprehensive and well-rounded framework for assessing human rights impacts.

Acknowledgments

This work was supported by the research grant ”Center for Digital CompliancE (DICE)” (VIL57420) from VILLUM FONDEN

References

[1] M. P. Brito, M. Stevenson, C. Bravo, Subjective machines: Probabilistic risk assessment based on deep learning of soft information, Risk Analysis 43 (2023) 516–529.
[2] Article 29 Working Party, WP29 Statement 14/EN WP 218 on the Role of a Risk-Based Approach to Data Protection Legal Frameworks, 2014. Adopted on 30 May 2014.
[3] H. Udsen, Databeskyttelsesret, 2022. 290-321.
[4] Article 29 Working Party, WP29 Guidelines on Data Protection Impact Assessment (DPIA), 2017. As last Revised and Adopted on 4 October 2017.
[5] International Organization for Standardization, ISO 31000:2009, Risk management — Principles and guidelines, 2009.
[6] International Organization for Standardization, ISO/IEC 29134, Information technology – Security techniques – Privacy impact assessment – Guidelines, 2017.
[7] International Organization for Standardization, Risk Management – Vocabulary – Guidelines for Use in Standards, Guide 73, 2002. Definition 3.1.1.
[8] S. Kogan, D. Levin, B. R. Routledge, J. S. Sagi, N. A. Smith, Predicting risk from financial reports with regression, in: Proceedings of the Human Language Technologies: The Annual Conference

of the North American Chapter of the Association for Computational Linguistics, Association for Computational Linguistics, Stroudsburg, PA, 2009, pp. 272–280.
[9] M. Surdeanu, R. Nallapati, G. Gregory, J. H. Walker, C. D. Manning, Risk analysis for intellectual property litigation, in: Proceedings of the 13th International Conference on Artificial Intelligence and Law, Pittsburgh, PA, USA, 2011.
[10] J. L. Leidner, F. Schilder, Hunting for the black swan: Risk mining from text, in: Proceedings of the ACL 2010 System Demonstrations, 2010, pp. 54–59.
[11] M. F. Tsai, C. J. Wang, On the risk prediction and analysis of soft information in finance reports, European Journal of Operational Research 257 (2017) 243–250.
[12] R. Clarke, Privacy impact assessment: Its origins and development, Computer Law & Security Review 25 (2009) 123–135.
[13] A. Mantelero, The fundamental rights impact assessment (fria) in the ai act: Roots, legal obligations and key elements for a model template, Computer Law & Security Review 54 (2024) 106020.
[14] Danish Institute for Human Rights (DIHR), Human Rights Impact Assessment Guidance and Toolbox, 2020.
[15] Commission Nationale de l'Informatique et des Libertés (CNIL), PIA Methodology: How to Carry out a Data Protection Impact Assessment, 2021. URL: https://www.cnil.fr/sites/cnil/files/atoms/files/cnil-pia-1-en-methodology.pdf.
[16] Agencia Española de Protección de Datos (AGPD), Guía para una Evaluación de Impacto en la Protección de Datos Personales (EIPD), 2014. Guidelines for a Personal Data Protection Impact Assessment.
[17] N. Götzmann (Ed.), Handbook on Human Rights Impact Assessment, Edward Elgar Publishing, 2019.
[18] A. Mantelero, Beyond Data: Human Rights, Ethical and Social Impact Assessment in AI, Springer Nature, 2022.