

Data Governance, Quality, and Accountability

Unpacking Article 10 of the EU AI Act

Nikoletta Constantinou

Attorney at Law, Special Teaching Staff, NUP School of Law
Researcher of the Jean Monnet Centre of Excellence AI-2-TRACE-CRIME

Georgios Pavlidis

UNESCO Chair & Jean Monnet Chair
Director of the Jean Monnet Center of Excellence AI-2-TRACE-CRIME, NUP
Associate Professor of International and EU Law School of Law, Neapolis University Pafos

Article 10 of the EU Artificial Intelligence Act introduces important requirements on data quality, governance, and traceability for high-risk AI systems. While the provision appears technical in nature, its rationale is fundamentally rights-based. Poor data practices, such as bias, incompleteness, or lack of documentation, can lead to discriminatory or opaque outcomes that directly affect individuals' fundamental rights. Article 10 seeks to mitigate these risks by establishing safeguards at the data level, from training to deployment. This study provides a close reading of the legal obligations set out in Article 10, situating them within the broader framework of the Charter of Fundamental Rights of the EU and the General Data Protection Regulation. Although Article 10 constitutes an important development in regulating the data foundations of AI, its effectiveness will ultimately depend on how its open-ended terms are interpreted and whether regulators are equipped to ensure accountability. This study was prepared in the context of the Jean Monnet Center of Excellence AI-2-TRACE-CRIME and was funded by the European Union. Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them.

1. Introduction

The EU Artificial Intelligence Act (AI Act),¹ adopted in 2024, is the first comprehensive legislative instrument of its kind and it marks a turning point in the regulation of AI technologies. The AI Act sets out a risk-based framework that imposes requirements on AI systems deemed to pose higher risks to health, safety, and fundamental rights. Among these requirements, Article 10 plays an important role, since it lays down detailed obligations regarding the quality and governance of data used in the development of high-risk AI systems. This is a domain where technical design choices are closely related to legal and ethical outcomes.

The connection between data and fundamental rights is neither incidental nor abstract. The training and testing of AI systems on biased, unrepresentative, or poor-quality data

can lead to discriminatory or erroneous outputs.² Individuals may not be aware of how these results came about or how they might challenge them. In fields such as employment, education, healthcare, or credit scoring, the consequences can be significant. Article 10 addresses the conditions under which data is collected, processed, and maintained, to prevent harms at their source. It is not merely about technical reliability, but about ensuring that automated systems respect the rights of the people they affect.

This study examines how Article 10 operationalises safeguards for fundamental rights through concrete obligations on data relevance, representativeness, error minimisation, traceability, and governance. It also considers how these obligations fit within the broader structure of the AI Act, and how they interact with existing EU legal frameworks, particularly the Charter of Fundamental Rights and the

1. 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), OJ L, 2024/1689, 12.7.2024,

2. Ferrara, Emilio, "The butterfly effect in artificial intelligence systems: Implications for AI bias and fairness" (2024) *Machine Learning with Applications* 15: 100525.

General Data Protection Regulation (GDPR).³ While Article 10 reflects the EU's commitment to trustworthy AI, it also raises questions about implementation, enforceability, and the extent to which these legal standards can be meaningfully translated into practice.

The discussion proceeds as follows. Section 2 examines the risks to fundamental rights posed by flawed or unregulated data practices in AI. Section 3 provides a close reading of Article 10, breaking down its legal components and exploring their normative significance. Section 4 situates Article 10 within the broader EU legal landscape, including overlaps with data protection. The final two sections present a critical analysis of whether Article 10 effectively fulfils its aim of embedding fundamental rights protections into the data practices of high-risk AI systems.

2. The Rights Risks of Data Failures in AI Systems

Much of the concern surrounding AI and its use in high-stakes areas (healthcare, social welfare, policing, etc.) centres on the quality and governance of the data that fuels these systems. While discussions of AI ethics and regulation often focus on algorithmic transparency or explainability,⁴ it is the data used to train, validate, and operate these systems that frequently determines whether their outcomes are fair, accurate, and lawful. In other words, data is not neutral. Data may reflect social patterns, institutional biases, and historical inequalities. Thus, when left unchecked, these patterns can be reproduced and reinforced by AI systems. This may result in negative consequences for individual rights and for social justice.

The risks are far from hypothetical. A growing body of evidence shows that biased or unrepresentative datasets have led to AI systems producing discriminatory or exclusionary outcomes, particularly for individuals from marginalised groups. For example, facial recognition technologies, one of the most academically studied areas within computer vision,⁵ have performed less accurately on women and people of colour, raising concerns under Article 21 of the EU Charter of Fundamental Rights, which prohibits discrimination on

any ground. Similarly, algorithmic risk assessment tools in criminal justice have been shown to embed racial bias, often drawing on proxy variables that reflect structural disadvantage.⁶ In such cases, there is a risk that poor data governance translates into violations of rights to equality, dignity, and due process.

Due to the complexity and opacity of many AI systems, individuals affected by automated decisions often cannot see how the underlying data was selected, labelled, or processed—and therefore cannot easily challenge outcomes they perceive as unfair. The lack of visibility into data governance practices may undermine the right to good administration (Article 41 of the Charter) and the right to an effective remedy (Article 47). Thus, the use of poor-quality or untraceable data in high-risk systems undermines the EU's commitment to protecting fundamental rights in the digital age.

This is the normative backdrop against which Article 10 of the AI Act must be understood. Its provisions are not aimed at technical optimisation for its own sake, but at ensuring that high-risk AI systems do not operate on problematic foundations and do not produce discriminatory, erroneous, or opaque outcomes. Datasets must be relevant, representative, and as free of errors as possible. Governance obligations around documentation and traceability can prevent harm before it occurs. Following this logic, Article 10 makes a shift from *ex post* redress to *ex ante* prevention.

3. Article 10 of the AI Act: Legal Architecture and Rights Functions

Article 10 of the EU AI Act lays out a detailed framework on data governance for high-risk AI systems. While the language is technical, its objectives are rooted in legal principles, in particular safeguarding the rights of individuals. More specifically, Article 10 seeks to ensure that the data underpinning automated decisions is not flawed, opaque, or unaccountable. In this sense, it translates rights-based concerns into operational requirements for AI developers and providers.

3.1. Ensuring Data Quality

The third paragraph of Article 10 requires that training, validation, and testing data of AI systems be relevant, representative, free of errors and complete, while also accounting for the specific characteristics or intended purpose of the system. On the surface, this reads like a list of quality indicators. However, each of these notions (relevance,

3. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) *OJ L 119*, 4.5.2016, pp. 1–88

4. Pavlidis, Georgios, “Unlocking the black box: Analysing the EU artificial intelligence act's framework for explainability in AI” (2024) *Law, Innovation and Technology*, Vol. 16.1, pp. 293–308; Balasubramaniam, Nagadivya, et al. “Transparency and explainability of AI systems: From ethical guidelines to requirements” (2023) *Information and Software Technology* 159: 107197.

5. Yucer, Seyma, et al., “Racial bias within face recognition: A survey” (2024) *ACM Computing Surveys*, Vol. 57.4, pp. 1–39.

6. O'Brien, Tim, “Compounding injustice: The cascading effect of algorithmic bias in risk assessments” (2021) *Geo. J. & Mod. Critical Race Persp.*, Vol. 13, p. 39 ff.

representativeness, error minimisation, completeness) is associated with legal and practical complexity.

The requirement of relevance is intended to ensure that only data necessary for the stated purpose of the AI system is used. This echoes the principle of purpose limitation found in Article 5(1)(b) of the GDPR.⁷ For its part, the notion of representativeness touches on questions of fairness and non-discrimination. If a dataset systematically underrepresents certain groups, the system may produce skewed results that disproportionately affect individuals or groups on the basis of race, gender, or other protected characteristics. This engages Article 21 of the Charter.⁸ Moreover, the obligation to use data free of errors, to the extent technically feasible, is particularly significant in systems that generate outcomes with legal or factual consequences, such as eligibility determinations or risk assessments. Finally, the notion of completeness seeks to guard against fragmented or partial data being used to train systems with potentially wide-ranging impact.

Yet the Act does not define these terms in detail, which leaves room for interpretation and, therefore, potential inconsistencies. Providers of AI systems will need to make judgments about what counts as “representative” or “complete,” often in the absence of clear benchmarks. This creates uncertainty, but also flexibility. If interpreted through a rights-based lens, these open terms can serve as guidelines for the scrutiny of data governance, especially in sectors where the risk of harm is acute.

In addition to this, Article 10 paragraph 4 of the AI Act requires that datasets reflect the specific context in which a high-risk AI system will operate, taking into account relevant geographical, contextual, behavioural, or functional characteristics. This obligation reinforces the principle of fitness-for-purpose.⁹ It ensures that AI systems are trained and validated with data relevant to their intended environment. This can improve accuracy of data. It also ties into the risk-based approach of the Act, demanding that providers anticipate variations in system performance across different settings to mitigate unintended harms or discriminatory outcomes.

3.2. A Framework for Data Governance

Article 10 paragraph 2 of the AI Act sets out a specific framework for the governance and management of datasets

used in the training, validation, and testing of high-risk AI systems.¹⁰ It requires that data practices be appropriate for the system’s intended purpose, reflecting the Act’s risk-based and proportionate approach. The provision covers key elements such as design choices, data provenance, and processing operations like annotation, cleaning, and aggregation, while also mandating the documentation of assumptions and an assessment of the adequacy, quality, and suitability of datasets. These obligations ensure that the development process is transparent, traceable, and aligned with related frameworks, most notably the GDPR, thereby supporting lawful, accountable, and trustworthy AI systems.

Points (f)–(h) of paragraph 2 highlight the critical focus on bias management and continuous improvement. Providers must assess datasets for potential biases that could harm health and safety, infringe fundamental rights, or lead to unlawful discrimination. They must also implement effective measures to detect, prevent, and mitigate such risks throughout the lifecycle of AI systems. They must identify and address data gaps or shortcomings that hinder compliance, adapting system design or sourcing additional data where necessary. Adherence to these practices enhances legal compliance and the reliability of high-risk AI systems, although implementation may pose practical challenges, especially for SMEs.

3.3. Processing Special Categories of Personal Data

Article 10 paragraph 5 of the AI Act establishes a strict and exceptional regime for processing special categories of personal data (data revealing racial or ethnic origin, political opinions, or health information), when it is strictly necessary to detect and correct bias in high-risk AI systems.¹¹ This provision recognises that certain biases cannot be effectively identified or mitigated without processing sensitive data.¹² However, it also reinforces the principle of data minimisation by mandating that such processing occur only when no alternative, including synthetic or anonymised data, is viable. The framework operates in conjunction with existing EU data protection rules under the GDPR, Regulation 2018/1725¹³ and

7. Forgó, Nikolaus, Stefanie Hähold, and Benjamin Schütze, “The principle of purpose limitation and big data” (2017) *New technology, big data and the law* (Singapore: Springer) pp. 17–42.
8. Iamicieli, Paola, et al., “EU Fundamental Rights and Non-Discrimination: Effective Protection in the Light of Article 21 of the Charter” (2022) *FRICoRE Casebook*, pp. 1–195.
9. Thompson, M., Fearn, T., “What exactly is fitness for purpose in analytical measurement?” (1996) *Analyst*, Vol. 121.3, pp. 275–278.

10. Janssen, Marijn, et al., “Data governance: Organizing data for trustworthy Artificial Intelligence” (2020) *Government information quarterly* 37: 101493.
11. Van Bekkum, Marvin, and Frederik Zuiderveen Borgesius, “Using sensitive data to prevent discrimination by artificial intelligence: Does the GDPR need a new exception?” (2023) *Computer Law & Security Review* 48: 105770.
12. Žliobaitė, Indrė, and Bart Custers. “Using sensitive personal data may be necessary for avoiding discrimination in data-driven decision models.” *Artificial Intelligence and Law* 24.2 (2016): 183–201.
13. Regulation (EU) 2018/1725 of the European Parliament and of the Council of 23 October 2018 on the protection of natural persons with regard to the processing of personal data by the Union institutions, bodies, offices and agencies and on the free

Directive 2016/680,¹⁴ ensuring legal coherence and protection of fundamental rights.

For the exceptional regime under Article 10 paragraph 5, the provision imposes stringent technical and organisational safeguards. Sensitive data must be pseudonymised and protected by state-of-the-art security measures, with strict access controls, documentation of processing activities, and obligations of confidentiality for authorised personnel. The data cannot be transferred or shared with third parties, and it must be deleted as soon as the bias is corrected or the retention period expires. Providers of AI systems must also maintain detailed records justifying why the processing of special categories of personal data was indispensable and why less intrusive methods could not achieve the same result. Therefore, a high level of accountability is achieved, which reflects the legislator's intent to balance bias mitigation with the protection of fundamental rights.

4. Article 10 in the Broader EU Rights Framework

Although Article 10 follows the specific logic of the AI Act, its legal and practical implications cannot be fully understood in isolation. The obligations it sets out engage wider EU legal principles, particularly those enshrined in the Charter of Fundamental Rights of the European Union and the GDPR. These frameworks provide interpretative context for understanding what Article 10 requires and what is at stake when it is not complied with.

From its opening articles and recitals, the AI Act stresses the need to ensure that AI systems do not undermine dignity, equality, or justice. Article 10 plays a central role in this effort: it seeks to embed rights protection at the level of the data, which is the very foundation of AI systems. The relevance of several Charter rights becomes apparent in this context. Article 21 of the Charter, which prohibits discrimination on grounds such as sex, race, ethnic origin, and disability, is directly implicated when AI systems reproduce or amplify historical bias encoded in data. Article 8, on the protection of personal data, also comes into play, particularly in the context of traceability and the lawful handling of training and testing datasets. The quality and completeness of data further impact Article 41, the right to good administration, and Article 47, which guarantees the right to an effective remedy and a fair hearing.

It is important to underline that these Charter rights are binding on both the EU institutions and Member States when implementing EU law.¹⁵ As such, national authorities tasked with enforcing the AI Act must ensure that interpretations of Article 10 are consistent with the provisions of the Charter. Therefore, there is a case for a rights-based reading of key terms in Article 10, which might otherwise be left to narrow technical interpretation.

The relationship between the AI Act and the GDPR is more complex.¹⁶ While the GDPR governs the processing of personal data and is rooted in the fundamental right to data protection, the AI Act introduces a layer of regulation focused specifically on system-level risk. Article 10 thus overlaps with certain GDPR provisions, raising questions about consistency and compliance.

One of the clearest points of intersection is the principle of data minimisation (Article 5(1)(c) GDPR),¹⁷ which may appear to conflict with Article 10's emphasis on completeness and representativeness. In practice, reconciling these principles will require careful design choices: developers must ensure that datasets include sufficient diversity to avoid bias, while not exceeding what is necessary for the system's intended function. Similarly, the GDPR's purpose limitation principle (Article 5(1)(b)) may constrain the reuse of datasets for different training objectives. This is particularly important when AI systems evolve after deployment.

Moreover, both instruments contain provisions on data subject rights, although with distinct focuses. The GDPR provides individual rights to access, rectification, and erasure of personal data. The AI Act, by contrast, operates at the level of system providers and imposes structural obligations that indirectly support individual rights, for example through traceability and documentation. While these mechanisms are complementary, the alignment between the two regimes is not seamless, especially in contexts where AI systems rely on large-scale datasets that are not easily attributable to identifiable individuals.

In this light, Article 10 can be seen as part of a broader EU regulatory architecture that treats data not simply as a resource to be managed, but as a legal object with implications for rights and accountability. The effectiveness of Article 10 will depend in part on how coherently it is

movement of such data, and repealing Regulation (EC) No 45/2001 and Decision No 1247/2002/EC, *OJL* 295, 21.11.2018, pp. 39–98

14. Directive (EU) 2016/680 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and on the free movement of such data, and repealing Council Framework Decision 2008/977/JHA, *OJL* 119, 4.5.2016, pp. 89–131

15. Lenaerts, Koen, "Exploring the Limits of the EU Charter of Fundamental Rights" (2012) *European Constitutional Law Review*, Vol. 8.3, pp. 375–403.

16. Musch, S., Borrelli, M., Kerrigan, C., "Bridging compliance and innovation: A comparative analysis of the EU AI Act and GDPR for enhanced organisational strategy" (2024) *Journal of Data Protection & Privacy*, Vol. 7.1, pp. 14–40.

17. Witt, C., De Bruyne, J., "The interplay between machine learning and data minimization under the GDPR: the case of Google's topics API" (2023) *International Data Privacy Law*, Vol. 13.4, pp. 284–298.

interpreted alongside the GDPR, and on the willingness of enforcement bodies to read both texts considering their shared underlying values.

5. Integration of ethical and societal critiques with Article 10's normative commitments.

While Article 10 represents a significant normative advance to tackle the risks of bias and discrimination, recurring concerns on transparency and accountability in algorithmic-decision making, call for an assessment within broader ethical implications. It could be argued that Article 10's effectiveness may be weakened by structural weaknesses in the overall regulatory framework and the lack of specificity and enforceability.¹⁸ Reconciling these tensions requires shifting from a product safety logic to a rights-based model of regulation¹⁹, with stronger provisions for independent oversight and civil participation.

Ethically, Article 10 represents a shift from ex post redress to ex ante prevention. Rather than waiting for harm to materialize and the offering remedies, the provision requires providers to proactively embed safeguards during the development phase. This anticipatory logic, utilizes abstract ethical principles into enforceable legal requirements such as transparency and accountability to succeed in the alignment of innovation and technological development with fundamental rights protection before deployment.

From an ethical standpoint, transparency has two main functions: empowerment and accountability. Empowerment allows individuals and interested parties to contest harmful or unfair outcomes, while accountability ensures that those who design and profit from AI systems get to answer for their consequences. According to Hohmann & Kollár, transparency is the one virtue that distinguishes and transforms AI from a "black box" technology²⁰ into a system open to democratic scrutiny, thus reinforcing both the rule of law and the people's trust. Yet, transparency provisions in the AI Act can be diluted by exemptions. As Hohmann & Kollár observe, high-risk AI systems must generally be registered in a public EU database (Article 71), but law enforcement, border management and migration authorities are only required to submit limited information to a non-public section of the database (Art. 49 par. 4). They conclude that in these contexts, advocacy groups and anyone acting in the public interest lack meaningful access

to the very systems most likely to infringe fundamental rights. This selective transparency undermines the ethical promise of Article 10, creating a double standard of accountability.

Article 10 advances accountability by creating ex ante obligations for providers that must demonstrate compliance with data quality and governance rules before deployment, traceability mechanisms and bias detection and correction responsibilities. This preventative model aligns with the ethical principle that those who design and profit from technology should also bear responsibility for its risks in contrast with older paradigms which leave the burden of seeking redress on individuals who may lack resources or access to justice.

Nevertheless, there is criticism about the lack of coherence in this accountability framework. Hohmann & Kollár took into account the work of the European Centre for Not-for-Profit Law and concluded that risk assessments rely primarily on self-assessment by providers, creating the danger of formalistic compliance rather than substantive safeguards. Similar observations were made for the inclusion of Fundamental Rights Impact Assessments (FRIAs). In effect, these practices risk promoting a compliance focused mindset, rather than a culture of genuine responsibility, ethical practices and continuous improvement, undermining the spirit of accountability.

Consequently, even with the combined effect of Article 10, the AI Act continues to attract criticism for adhering too closely towards a product safety logic, treating AI systems as if they were consumer goods to be certified and placed on the market²¹, despite its bona fide attempt to foster an environment of accountability oriented towards a rights-based approach.

Article 10, sits awkwardly between these two logics and seeks to navigate the ethical and legal landscape surrounding AI, fundamental rights and data protection. While it adopts a preventative approach aligned with ethical considerations, its reliance on technical standards may reduce substantive rights protections to mere bureaucratic formalities. To reconcile the ethical intent of Article 10 with ongoing criticisms, it is important to address the inherent tensions within AI, fundamental rights and data protection laws. This includes ensuring that ethical commitments are integrated into the regulatory framework, promoting genuine rights protection, rather than superficial compliance.

6. Concluding Remarks

Article 10 of the AI Act sets out the ambition to ensure high-risk AI systems are built on reliable, well-governed, and traceable data. However, the practical and legal challenges

18. Chiappetta, A., "Navigating the AI frontier: European parliamentary insights on bias and regulation, preceding the AI Act" (2023) *Internet Policy Review*, Vol. 12(4), pp. 1–26.

19. Hohmann, B., & Kollár, G., "Reflections on the data protection compliance of AI systems under the EU AI Act" (2025) *Cogent Social Sciences*, Vol. 11(1). <https://doi.org/10.1080/23311886.2025.2560654>

20. Latour, B., *Pandora's Hope: Essays on the Reality of Science Studies* (Harvard University Press, 1999).

21. Svitych, O., "Blind transparency: A critical discourse analysis of the EU AI Act" (2025) *Critical Policy Studies*, Latest Articles, pp. 1–17.

involved in turning this ambition into reality are considerable. As with many complex regulatory instruments, much will depend on interpretation, implementation, and enforcement.

One of the most immediate difficulties has to do with the indeterminate nature of key concepts used in Article 10. Notions such as “representative,” “free of errors,” or “appropriate data governance practices” leave significant discretion to providers, notified bodies, and national competent authorities to interpret obligations. This may lead to divergent practices across EU Member States. In highly sensitive contexts, such as credit scoring or migration control, what one actor considers “representative” might be deeply contested by civil society or affected communities.

Moreover, the requirement that data be as free of errors as possible introduces a degree of flexibility, but also a potential loophole. How much error is too much? What level of inaccuracy is acceptable before an AI system is deemed non-compliant? If there are no technical benchmarks or risk thresholds or harmonised guidelines tied explicitly to rights protection, enforcement bodies may struggle to assess whether a provider has met its obligations under Article 10.

These uncertainties are not purely academic. They create real risks of inconsistent implementation and regulatory arbitrage. They may also lead to overly cautious enforcement. Smaller providers, especially SMEs, may lack the technical and legal capacity to navigate these grey zones, while larger actors may be tempted to adopt minimalist interpretations unless supervised effectively.

Supervisory authorities and market surveillance bodies will bear much of the burden in giving life to Article 10’s safeguards. Their ability to conduct audits, demand documentation, and impose penalties will be essential. But without sufficient resources, expertise, and a clear mandate to prioritise rights protection over formal compliance, these bodies may struggle to enforce Article 10, especially when faced with complex, opaque, or proprietary systems.

Finally, Article 10 must be understood not as an isolated compliance burden, but as a mechanism for embedding human rights into the lifecycle of AI systems. If interpreted and applied with this spirit in mind, it could serve as a model for data governance in algorithmic regulation more broadly—not just in the EU, but globally.

References

- Balasubramaniam, Nagadivya, et al. “Transparency and explainability of AI systems: From ethical guidelines to requirements” (2023) *Information and Software Technology* 159: 107197.
- Chiappetta, A., “Navigating the AI frontier: European parliamentary insights on bias and regulation, preceding the AI Act” (2023) *Internet Policy Review*, Vol. 12(4), pp. 1–26.
- Ferrara, Emilio, “The butterfly effect in artificial intelligence systems: Implications for AI bias and fairness” (2024) *Machine Learning with Applications* 15: 100525.
- Forgó, Nikolaus, Stefanie Händel, and Benjamin Schütze, “The principle of purpose limitation and big data” (2017) *New technology, big data and the law* (Singapore: Springer) pp. 17–42.
- Hohmann, B., & Kollár, G., “Reflections on the data protection compliance of AI systems under the EU AI Act” (2025) *Cogent Social Sciences*, Vol. 11(1). <https://doi.org/10.1080/23311886.2025.2560654>
- Iamiceli, Paola, et al., “EU Fundamental Rights and Non-Discrimination: Effective Protection in the Light of Article 21 of the Charter” (2022) *FRICoRE Casebook*, pp. 1–195.
- Janssen, Marijn, et al., “Data governance: Organizing data for trustworthy Artificial Intelligence” (2020) *Government information quarterly* 37: 101493.
- Latour, B., *Pandora’s Hope: Essays on the Reality of Science Studies* (Harvard University Press, 1999).
- Lenaerts, Koen, “Exploring the Limits of the EU Charter of Fundamental Rights” (2012) *European Constitutional Law Review*, Vol. 8.3, pp. 375–403.
- Musch, S., Borrelli, M., Kerrigan, C., “Bridging compliance and innovation: A comparative analysis of the EU AI Act and GDPR for enhanced organisational strategy” (2024) *Journal of Data Protection & Privacy*, Vol. 7.1, pp. 14–40.
- O’Brien, Tim, “Compounding injustice: The cascading effect of algorithmic bias in risk assessments” (2021) *Geo. J.L. & Mod. Critical Race Persp.*, Vol. 13, p. 39 ff.
- Pavlidis, Georgios, “Unlocking the black box: Analysing the EU artificial intelligence act’s framework for explainability in AI” (2024) *Law, Innovation and Technology*, Vol. 16.1, pp. 293–308, <https://doi.org/10.1080/17579961.2024.2313795>
- Svitych, O., “Blind transparency: A critical discourse analysis of the EU AI Act” (2025) *Critical Policy Studies*, Latest Articles, pp. 1–17. <https://doi.org/10.1080/19460171.2025.2496193>
- Thompson, M., Fearn, T., “What exactly is fitness for purpose in analytical measurement?” (1996) *Analyst*, Vol. 121.3, pp. 275–278.
- Van Bekkum, Marvin, and Frederik Zuiderveen Borgesius, “Using sensitive data to prevent discrimination by artificial intelligence: Does the GDPR need a new exception?” (2023) *Computer Law & Security Review* 48: 105770.
- Witt, C., De Bruyne, J., “The interplay between machine learning and data minimization under the GDPR: the case of Google’s topics API” (2023) *International Data Privacy Law*, Vol. 13.4, pp. 284–298.
- Yucer, Seyma, et al., “Racial bias within face recognition: A survey” (2024) *ACM Computing Surveys*, Vol. 57.4, pp. 1–39.