

Right-to-Explanation Principles in Automated Public Administration

Dr. Alexander Bennett

Professor, University of Edinburgh, Edinburgh, United Kingdom

Abstract

Automated and artificial-intelligence-supported systems are increasingly used within public administration to assess eligibility for benefits, prioritize inspections, evaluate applications, detect suspected irregularities, support migration and border decisions, allocate public services, and structure other administrative processes. Although automation may improve consistency, processing capacity, and administrative efficiency, it can weaken procedural accountability when affected individuals cannot understand why a decision was reached or how algorithmic outputs influenced official reasoning. This article develops a rights-oriented framework for explanation in automated public administration.

It distinguishes technical model explainability from legally meaningful explanation and argues that an adequate administrative explanation should identify the role of automation, the material information used, the principal factors affecting the individual outcome, the relationship between those factors and the final decision, significant limitations, and a practical route for correction, human review, or contestation. Contemporary European law provides particularly important developments. The General Data Protection Regulation requires meaningful information concerning the logic, significance, and envisaged consequences of qualifying automated decision-making, while the Court of Justice of the European Union strengthened the practical content of these rights in *SCHUFA Holding* and, particularly, *Dun & Bradstreet Austria*.

The latter judgment confirmed in February 2025 that meaningful information must enable the individual to understand how personal data were actually used to obtain the particular automated result. The EU Artificial Intelligence Act separately contains an express Article 86 right to obtain clear and meaningful explanations concerning certain individual decisions based on outputs of Annex III high-risk AI systems. A synthetic analytical framework comprising 480 hypothetical administrative decisions compares four explanation conditions: no substantive explanation, generic notification of AI use, case-specific explanation of material reasons, and a comprehensive model combining individualized reasons with contestability safeguards.

The composite procedural-accountability score increases from 34.9 to 51.4, 73.5, and 86.6 respectively. These values are illustrative rather than empirical. The study concludes that a meaningful right to explanation should operate as a gateway to effective administrative justice rather than as a narrow requirement to describe algorithmic technology. Explanation becomes legally significant when it enables an affected person to understand the decision, identify a potentially material error, and exercise available rights of review or challenge.



Keywords: right to explanation, automated public administration, administrative law, artificial intelligence, automated decision-making, algorithmic accountability, procedural fairness, GDPR, EU AI Act, contestability

1. Introduction

Public administration is undergoing a substantial transformation as government agencies integrate automated systems into activities that were previously performed entirely through human administrative judgment. Algorithmic systems may screen applications, estimate risk, rank cases for investigation, identify suspected fraud, recommend eligibility decisions, determine priority for public resources, or supply predictive information to administrative officials. Such tools can increase administrative capacity and may support more consistent processing of large caseloads. Their use nevertheless becomes legally sensitive when automated outputs affect access to social benefits, employment-related public services, healthcare, education, immigration processes, law enforcement, licensing, taxation, or other legally significant interests. The EU Artificial Intelligence Act itself identifies several areas involving access to essential public services and benefits, law enforcement, migration, administration of justice, education, and related activities as potentially high-risk contexts.

The principal challenge is not simply whether an algorithm is technically accurate. Public administration is governed by principles that require governmental power to remain accountable, reviewable, and compatible with fundamental rights. An individual adversely affected by an administrative decision ordinarily needs more than notification that an unfavorable outcome occurred. Effective challenge depends on understanding the factual and legal basis of the decision sufficiently to determine whether relevant information was incorrect, an impermissible factor was considered, an important circumstance was omitted, or administrative discretion was exercised improperly. Automation can disrupt this traditional relationship between decision, reasons, and review when institutional actors themselves rely heavily on scores, classifications, or predictions produced by complex systems. Cobbe's analysis of automated public-sector decision-making demonstrates why administrative law, data-protection law, and technical understanding must be considered together when assessing algorithmically supported government decisions.

The phrase right to explanation has historically generated considerable legal debate. Wachter, Mittelstadt, and Floridi argued in 2017 that the GDPR did not contain the broad ex-post right to explanation frequently attributed to it, whereas Selbst and Powles argued that the GDPR's requirement to provide meaningful information should be read more substantively. Subsequent judicial developments have significantly changed the practical landscape. In its February 27, 2025 judgment in *Dun & Bradstreet Austria*, the Court of Justice held that meaningful information under GDPR Article 15(1)(h) must cover the procedure and principles actually applied so that an affected person can understand which personal data were used and how they contributed to the automated result. Recent scholarship therefore describes the earlier debate over whether some explanatory entitlement exists under the GDPR as substantially overtaken by CJEU case law, while emphasizing that GDPR information rights and the AI Act's explicit Article 86 right remain legally distinct.

The EU Artificial Intelligence Act introduces another important layer. Article 86 is expressly titled “Right to explanation of individual decision-making.” It applies to an affected person subject to a



qualifying decision taken by a deployer on the basis of the output of specified Annex III high-risk AI systems where that decision produces legal effects or comparably significant effects involving adverse impacts on health, safety, or fundamental rights. The person is entitled to obtain clear and meaningful explanations concerning the role played by the AI system and the main elements of the decision. The EU implementation timeline must nevertheless be stated accurately: following the Digital Omnibus changes that entered into force on July 27, 2026, the main rules governing Annex III high-risk systems are scheduled to apply from December 2, 2027, while rules concerning high-risk AI embedded in covered physical products are scheduled from August 2, 2028.

This study therefore examines explanation not as a purely technical property of an AI model but as a procedural requirement connecting administrative power with accountability. It proposes that a legally useful explanation should help an affected person answer four practical questions: Was automation materially involved? Why did this particular outcome occur? What information or assumptions could materially change the outcome? How can the decision be corrected or challenged? The study combines legal analysis with a transparent simulation to demonstrate how different levels of explanation could be evaluated in future empirical research. It does not claim that one universal explanation format is legally required across all jurisdictions; instead, it develops a principled framework capable of adaptation to different public-law systems.

2. Objectives of the Study

2.1 To Clarify the Legal Meaning of Explanation in Automated Administration

The first objective is to distinguish meaningful legal explanation from technical descriptions of an AI model. The analysis focuses on the relationship among transparency, individualized reasons, data-use information, procedural fairness, and administrative review.

2.2 To Examine the Relationship Between Explanation and Contestability

The second objective is to analyze whether explanation becomes practically effective only when it enables an affected person to identify relevant errors, request human reconsideration, express an opposing view, or use an administrative or judicial remedy.

2.3 To Develop an Operational Explanation Framework

The third objective is to develop an evaluation structure for public authorities using automated or AI-supported decision systems. The proposed structure incorporates disclosure of AI involvement, individualized reasoning, material factors, data traceability, intelligibility, review pathways, and documentation.

3. Review of Literature

3.1 From Algorithmic Transparency to Meaningful Explanation

Early literature on algorithmic accountability frequently concentrated on opacity. Complex computational systems may produce outputs whose internal operations are difficult for ordinary citizens, public officials, and sometimes even system developers to interpret. Yet disclosure of mathematical

architecture is rarely sufficient for administrative justice. Technical transparency and legal explanation serve different functions.

Wachter, Mittelstadt, and Floridi's influential 2017 analysis questioned whether the GDPR created the extensive ex-post right to explanation popularly attributed to it. Their position stimulated extensive debate about the distinction between information about system functionality and explanations of a specific decision.

Selbst and Powles responded that the requirement to provide “meaningful information” should be interpreted in light of the GDPR's broader transparency and safeguards architecture. Their approach emphasized that information must be sufficiently useful to protect the individual's ability to understand and challenge automated decision-making.

The later CJEU jurisprudence significantly strengthens this functional approach. In *SCHUFA Holding*, the Court interpreted Article 22 in a manner intended to prevent organizations from avoiding automated-decision safeguards merely because a score is formally produced by one entity and relied upon by another. The Court emphasized protection against risks arising from decisions effectively determined through automated processing.

The February 2025 *Dun & Bradstreet Austria* judgment provides greater specificity. The Court held that the GDPR requirement for meaningful information encompasses relevant information concerning the procedures and principles used to process personal data for a particular automated result. The explanation must be concise, transparent, intelligible, and easily accessible. Importantly, model complexity does not relieve a controller of the obligation to explain. The Court also indicated that, in a scoring context, information showing how changes in relevant personal data would have produced a different result may provide an intelligible form of explanation.

This judicial development supports an important distinction: meaningful explanation need not expose every mathematical calculation or source-code instruction. Its purpose is to make the particular decision comprehensible enough for the affected person to exercise legal rights.

3.2 Explanation as an Administrative-Law Requirement

Automated administration raises issues extending beyond data protection. Administrative law has long required public power to remain subject to legality, rationality, procedural fairness, and review. Automation does not necessarily remove administrative discretion; rather, it can redistribute discretion across software design, data classification, system procurement, configuration, threshold setting, and the later interpretation of algorithmic outputs.

McCann argues that automated systems can reshape relationships among traditional administrative actors and potentially constrain human discretion rather than becoming autonomous sites of legal discretion themselves. Covilla similarly argues that AI-supported administrative discretion must remain bounded by rule-of-law requirements, particularly duties of care, reason-giving, and judicial review.

Mir's 2025 examination of the AI Act from an administrative-law perspective emphasizes the relationship between Article 86 and established duties to provide reasons for administrative decisions. In public administration, therefore, an explanation cannot be reduced to telling a citizen that an AI system

contributed 37% to an internal score. The explanation needs to connect computational output to the legal and factual reasoning by which the public authority reached its decision.

This distinction is especially important when AI is used only as decision support. A public authority may argue that the algorithm did not formally make the decision because an official signed or approved it. Yet if the official systematically adopts the system's recommendation, the citizen's procedural position may depend substantially on understanding the algorithmically produced information.

3.3 Explanation, Human Rights, and Effective Remedy

Explanation also functions as an enabling condition for remedies. The Council of Europe's Framework Convention on Artificial Intelligence requires context-sensitive transparency and oversight measures and establishes accountability and responsibility principles for adverse impacts associated with AI systems. More specifically, Article 14 of the Convention addresses remedies. It requires measures designed to ensure that relevant information concerning AI systems capable of significantly affecting human rights is documented and, where appropriate, made available or communicated to affected persons. The information should be sufficient to allow affected persons to contest decisions made or substantially informed by AI, while an effective opportunity to lodge a complaint should also be available. Article 15 separately requires effective procedural safeguards where AI systems significantly affect the enjoyment of human rights.

These provisions reveal why explanation and contestation should not be treated as independent concepts. A highly detailed explanation without an available procedure for correction may have limited legal value. Conversely, a nominal appeal mechanism cannot function effectively if the person does not know what factual or automated reasoning needs to be challenged.

The literature therefore supports a broader conception of the right to explanation as part of an administrative accountability chain consisting of notice, intelligibility, individualized reasons, correction, human reconsideration, record preservation, and effective remedy.

4. Research Methodology

4.1 Research Design and Legal Analytical Method

The study adopts a combined doctrinal and simulation-based methodological design. The doctrinal component examines legal principles emerging from the GDPR, EU AI Act, CJEU jurisprudence, the Council of Europe AI Framework Convention, and administrative-law scholarship.

The research does not claim that these sources create an identical right across every jurisdiction. Rather, they are used to identify recurring principles that can inform a general explanation framework for automated public administration.

The simulation component demonstrates how explanation architectures could be compared empirically. Four hypothetical explanation conditions were constructed, representing progressively stronger procedural safeguards.

Table 1: Explanation Conditions in the Simulation Framework

Explanation Condition	Simulated Decisions	Main Characteristics	Procedural Function
No substantive explanation	120	Citizen receives outcome with minimal reason information	Administrative baseline
Generic AI notice	120	Citizen is told AI or automation contributed to the process, but receives little individualized reasoning	Transparency of system involvement
Case-specific reasons	120	Explanation identifies material facts, principal decision factors, AI role, and relationship between factors and outcome	Individual understanding
Reasons + contestability	120	Case-specific explanation plus correction mechanism, meaningful human review, challenge pathway, and decision record	Effective procedural accountability

Note: All observations are synthetic and were created solely for methodological illustration.

4.2 Explanation Evaluation Criteria

Four domains were used to evaluate the simulated explanation conditions.

Decision comprehension measures whether the affected person can understand why the specific administrative outcome occurred.

Factor transparency measures whether the person can identify the principal personal data, factual considerations, thresholds, or algorithmically derived factors that materially influenced the decision.

Error-detection capacity represents the person's ability to identify incorrect information, misclassification, omitted circumstances, or a potentially inappropriate inference.

Contestability readiness represents whether the explanation provides enough information and procedural direction to support correction, reconsideration, appeal, or another available remedy.

Each variable was represented on a 0–100 scale. The composite procedural-accountability score was calculated as the arithmetic mean of the four dimensions.

The framework deliberately does not include disclosure of source code as a universal criterion. Contemporary legal developments indicate that an explanation can be meaningful without requiring disclosure of every technical operation. The CJEU has instead emphasized intelligibility and the practical ability to understand how relevant data were used in producing the result.

4.3 Simulation Procedure, Validity, and Ethics

A hypothetical analytical sample of 480 administrative decisions was modeled, with 120 observations allocated to each explanation condition. The underlying decision was assumed to concern access to a public benefit for which an automated risk or eligibility assessment contributed materially to the administrative process.

The substantive eligibility rules were assumed to remain constant across conditions. Only the amount and quality of information available to the affected individual changed.

The study does not present inferential significance testing because the values were deliberately generated for methodological demonstration. Statistical significance based on synthetic observations could create a misleading impression that an actual citizen experiment had been completed.

A future empirical study should recruit real participants, randomly expose them to standardized administrative notices, measure actual understanding and error-detection performance, and preregister hypotheses and outcome measures. Any real study involving administrative records or vulnerable benefit recipients would also require appropriate privacy, ethical, and institutional safeguards.

5. Analysis

5.1 Comparative Explanation Performance

The simulated results demonstrate a substantial difference between notifying a citizen that automation was used and explaining how automation contributed to the individual's outcome.

The no-explanation condition produced a composite procedural-accountability score of 34.9. A generic AI notice increased the score to 51.4, indicating improved awareness of automated involvement but only modest gains in error detection and contestability.

Table 2: Simulated Outcomes Across Explanation Conditions

Condition	Decision Comprehension Mean ± SD	Factor Transparency Mean ± SD	Error- Detection Capacity Mean ± SD	Contestability Readiness Mean ± SD	Composite Score
No substantive explanation	38.6 ± 10.4	34.9 ± 9.8	29.7 ± 10.6	36.4 ± 9.7	34.9
Generic AI notice	57.8 ± 9.3	51.2 ± 10.1	44.6 ± 9.4	52.0 ± 9.0	51.4
Case-specific reasons	77.4 ± 8.0	72.6 ± 8.7	69.3 ± 9.2	74.7 ± 8.4	73.5
Reasons + contestability	87.8 ± 7.2	84.1 ± 7.4	86.3 ± 7.9	88.2 ± 6.8	86.6

Source: Author's synthetic simulation. Values are illustrative and do not describe observed government decisions or actual citizens.

5.2 Explanation Must Be Individualized

The modeled increase from 51.4 under a generic AI notice to 73.5 under case-specific reasons captures an important legal distinction.

A notice stating that “an automated system was used to assist this decision” satisfies one form of transparency, but it does not tell the individual why the decision went against them.


A meaningful individualized explanation should generally enable the recipient to understand the connection among relevant facts, algorithmic processing, applicable criteria, and final administrative reasoning.

The *Dun & Bradstreet Austria* judgment strengthens this interpretation in the GDPR context. The CJEU held that meaningful information must enable the person to understand which personal data were used and how they were used to obtain the particular automated result. The Court rejected the proposition that computational complexity can itself excuse inadequate explanation.

The judgment also provides a useful practical insight: explaining how a material change in relevant input information could have changed the result may sometimes be more understandable than disclosing a mathematical formula.

5.3 Explanation and Procedural Accountability

Graph 1: Simulated Procedural Accountability Across Explanation Conditions



The graph demonstrates the modeled progression from formal transparency toward operational procedural justice. Generic disclosure improves the citizen's awareness that automation was used, but the greatest improvement occurs when the administrative body explains the individualized reasons underlying the outcome.

The final condition produces an additional increase because explanation is connected directly with correction and review.

This relationship reflects the broader human-rights structure of the Council of Europe AI Framework Convention, which links relevant AI information with the affected person's capacity to contest decisions and obtain access to effective procedural safeguards.

The key analytical implication is that the quality of explanation should be measured by what the affected person can do with it, not merely by how much technical information is disclosed.

6. Discussion

6.1 A Right to Explanation Is Not a Right to Source Code

One of the persistent difficulties in algorithmic-governance debates is the assumption that meaningful explanation requires complete disclosure of model architecture, source code, or proprietary algorithms. Such disclosure may occasionally be relevant to regulators, courts, auditors, or technical experts, but it should not be treated as the ordinary citizen-facing standard.

A highly technical explanation can be formally detailed yet practically unintelligible. The CJEU's 2025 interpretation is important precisely because it focuses on the individual's understanding of the procedure and principles actually applied to their personal data rather than demanding a complete technical reconstruction of every computational step.

Trade-secret or third-party confidentiality interests also do not necessarily eliminate explanatory obligations. *Dun & Bradstreet Austria* recognizes that competing interests may require balancing through competent authorities or courts rather than permitting a controller simply to refuse meaningful information altogether.

For public administration, an effective explanation should therefore be designed at two levels. The citizen-facing level should communicate the material reasons in understandable language. The audit-facing level should preserve deeper technical documentation, data lineage, system versions, relevant model information, and decision records for authorized oversight.

6.2 Seven Principles of a Meaningful Administrative Explanation

The legal and methodological analysis supports seven principles for public authorities.

First, disclosure of automated involvement. A person should know when an AI or automated system materially contributed to a consequential administrative decision. The EU AI Act already establishes information duties for certain high-risk systems designed to make or assist decisions about natural persons.

Second, decision-specific reasons. An explanation should address why this particular person received this particular outcome. Generic descriptions of how a model works are insufficient when individualized reasoning is necessary for challenge.

Third, material-factor transparency. The explanation should identify the main information or categories of information that materially affected the result and describe their relevance.

Fourth, intelligibility. Information should be presented in language and structure that an ordinary affected person can reasonably understand. This principle corresponds closely with the CJEU's requirement that GDPR information be concise, transparent, intelligible, and accessible.

Fifth, sensitivity to outcome-changing information. Where useful and legally appropriate, authorities should explain whether correction or alteration of a material input could have changed the decision. Counterfactual-style information can be particularly useful for identifying factual errors and understanding decisional thresholds.

Sixth, human reconsideration and contestability. Explanation should identify how the person can correct data, present additional circumstances, seek human reconsideration, or challenge the decision under the applicable administrative procedure. The Council of Europe Framework Convention directly connects sufficient AI-related information with the ability to contest decisions and lodge complaints.

Seventh, reproducible administrative records. The authority should preserve enough internal information to reconstruct the AI-assisted decision when a complaint, audit, administrative appeal, ombudsman investigation, or judicial review occurs. Without such records, an apparently clear citizen-facing explanation may not be verifiable.

6.3 Relationship Between the GDPR, AI Act, and Administrative Justice

Legal precision is important because the GDPR and AI Act should not be treated as interchangeable explanation regimes.

GDPR Articles 13, 14, and 15 contain information rights relating to qualifying automated decision-making and require meaningful information about the logic involved, significance, and envisaged consequences. Article 22 additionally addresses decisions based solely on automated processing that produce legal or similarly significant effects, subject to specified conditions and safeguards.

The AI Act takes a different regulatory approach. Article 86 expressly creates a right to obtain clear and meaningful explanations concerning the role of specified high-risk AI outputs and the main elements of qualifying decisions. Unlike Article 22 GDPR, Article 86 is not conceptually limited to a decision being solely automated; it addresses decisions made by a deployer on the basis of covered AI outputs.

The AI Act also places substantial obligations around high-risk systems, including transparency, human oversight, documentation, and—for specified public authorities and public-service deployers—fundamental-rights impact assessment. As of August 2026, however, researchers and administrators must account for the revised implementation calendar: the high-risk rules are scheduled to apply from December 2, 2027, with later application for covered product-embedded high-risk systems.

7. Conclusion

Automated public administration creates an apparent paradox. Government can process more information and reach decisions more rapidly, yet the same technological capacity can make individual decisions harder to understand and challenge. This difficulty is particularly serious because administrative decisions may determine access to benefits, services, legal status, employment opportunities, education, healthcare, or other interests connected with fundamental rights.

The central finding of this study is that a meaningful right to explanation should not be understood as a demand for technical disclosure alone. Explanation is valuable because it supports understanding, error detection, human reconsideration, and effective remedy.



Contemporary European legal developments reinforce this functional interpretation. The GDPR provides meaningful-information safeguards for qualifying automated decision-making; SCHUFA Holding strengthened protection against attempts to separate algorithmic scoring artificially from consequential downstream decisions; and *Dun & Bradstreet Austria* clarified that Article 15(1)(h) requires understandable information concerning the procedures and principles actually applied to the individual's data. The EU AI Act goes further terminologically by establishing an express Article 86 right to explanation for qualifying decisions based on specified high-risk AI-system outputs.

The Council of Europe Framework Convention broadens the governance perspective by connecting transparency, accountability, remedies, sufficient decisional information, complaint mechanisms, and procedural safeguards.

The synthetic analysis illustrates how these principles could be evaluated empirically. Generic notification that AI was used produces only moderate simulated gains in procedural accountability. Case-specific reasoning performs substantially better, while the strongest modeled condition combines individualized explanation with correction, review, and contestation.

These numerical results are not empirical evidence. Their purpose is to demonstrate that future studies should evaluate explanations according to functional outcomes rather than technical completeness.

For public agencies, the practical standard should therefore be straightforward: an affected person should be able to determine what happened, why it happened, what information materially influenced the outcome, whether a relevant error can be corrected, and where the decision can be challenged.

An explanation that cannot support these functions may satisfy transparency in form while failing administrative justice in substance.

References

1. Wachter, S., Mittelstadt, B. D., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99.
2. Wachter, S., Mittelstadt, B. D., & Russell, C. (2018). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2), 841–887.
3. Goodman, B., & Flaxman, S. (2017). European Union regulations on algorithmic decision-making and a “right to explanation.” *AI Magazine*, 38(3), 50–57.
4. Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). Accountable algorithms. *University of Pennsylvania Law Review*, 165(3), 633–705.
5. Burrell, J. (2016). How the machine “thinks”: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12.
6. Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press.
7. Barocas, S., & Selbst, A. D. (2016). Big data’s disparate impact. *California Law Review*, 104(3), 671–732.



8. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59–68.
9. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.
10. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), Article 93.
11. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
12. Diakopoulos, N. (2016). Accountability in algorithmic decision making. *Communications of the ACM*, 59(2), 56–62.
13. Edwards, L., & Veale, M. (2017). Slave to the algorithm? Why a “right to an explanation” is probably not the remedy you are looking for. *Duke Law & Technology Review*, 16, 18–84.
14. Wirtz, B. W., Weyerer, J. C., & Geyer, C. (2019). Artificial intelligence and the public sector—Applications and challenges. *International Journal of Public Administration*, 42(7), 596–615.
15. Janssen, M., Brous, P., Estevez, E., Barbosa, L. S., & Janowski, T. (2020). Data governance: Organizing data for trustworthy artificial intelligence. *Government Information Quarterly*, 37(3), 101493.
16. Valle-Cruz, D., Criado, J. I., Sandoval-Almazán, R., & Ruvalcaba-Gomez, E. A. (2020). Assessing the public policy-cycle framework in the age of artificial intelligence: From policy planning to policy evaluation. *Government Information Quarterly*, 37(4), 101509.
17. Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28, 689–707.
18. European Union. (2016). *Regulation (EU) 2016/679: General Data Protection Regulation*. Official Journal of the European Union. The regulation contains important provisions concerning automated individual decision-making, profiling, transparency, and safeguards.
19. Article 29 Data Protection Working Party. (2018). *Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679*. European Commission.
20. Organisation for Economic Co-operation and Development. (2019). *Recommendation of the Council on Artificial Intelligence*. OECD. The recommendation emphasizes trustworthy AI, transparency, explainability, accountability, and human-centered governance.