

Algorithmic Transparency Dashboards for Citizen-Facing Decision Systems

Dr. Olivia M. Grant

Associate Professor, University College Dublin, Ireland

Abstract

Algorithmically supported decision systems increasingly influence citizen interactions with public institutions, including eligibility assessment, case prioritization, inspection, licensing, taxation, social protection, fraud detection, and service routing. Transparency initiatives have consequently expanded from general policy commitments toward public algorithm registers, algorithmic impact assessments, standardized disclosure records, and increasingly interactive forms of public reporting. The United Kingdom's Algorithmic Transparency Recording Standard provides a standardized mechanism through which public bodies disclose how and why algorithmic tools are used, and its use became mandatory across central government in 2021. The Netherlands now maintains a national Algorithm Register containing more than 1,500 published algorithm descriptions, while Helsinki's AI Register combines system information with public feedback mechanisms.

This study develops a governance framework for algorithmic transparency dashboards designed specifically for citizen-facing decision systems. Unlike static disclosures that simply announce the presence of an algorithm, the proposed dashboard model combines purpose specification, decision-role disclosure, data provenance, performance information, fairness and bias indicators, human-oversight arrangements, institutional ownership, change history, and practical routes for explanation, correction, and appeal. A conceptual-methodological design is combined with a simulation-based comparison of five transparency models ranging from basic disclosure pages to citizen-centered dashboards. The simulated analysis indicates that greater disclosure completeness does not automatically generate citizen actionability; meaningful transparency requires information architecture that converts technical disclosure into understandable and usable civic knowledge. The citizen-centered dashboard model achieves the strongest illustrative combination of disclosure completeness, actionability, and accountability support. The study concludes that public-sector transparency should be evaluated not by how much technical information is released, but by whether affected individuals can discover that an algorithm is being used, understand its role, assess relevant risks, identify the responsible authority, and act when they believe a decision is incorrect or unfair.

Keywords: *algorithmic transparency, public-sector AI, transparency dashboard, algorithmic accountability, citizen-facing systems, explainable AI, algorithm register, public administration, contestability, digital government*



1. Introduction

Algorithmic systems have become increasingly embedded in the operational infrastructure of contemporary government. Public bodies use automated and semi-automated tools to classify applications, rank cases, detect anomalies, allocate investigative attention, forecast demand, assess risk, route requests, support benefit administration, and assist public officials in making administrative decisions. This expansion has generated substantial concern because citizens frequently experience the outcome of an algorithmically mediated process without knowing that an algorithm was involved, what information influenced the decision, whether humans reviewed the result, or how an error can be corrected. Levy, Chaselow, and Riley's survey of public-sector algorithmic decision-making emphasizes that government algorithms create important questions concerning accountability, privacy, inequality, procurement, participation, and evaluation across the full system lifecycle.

Transparency has therefore become a major component of public-sector algorithmic governance. NIST distinguishes transparency from explainability and interpretability: transparency helps answer what happened in a system, explainability addresses how a decision was produced, and interpretability concerns why a system output has meaning in the relevant context. This distinction is particularly useful for citizen-facing decision systems because disclosure of technical documentation alone may provide formal transparency without giving an affected person the information required to understand or challenge a decision. A citizen denied a permit or referred for additional investigation may gain little from knowing the name of a machine-learning architecture if the system does not also explain which administrative process it supports, what information was used, whether a human made the final determination, and where an appeal can be submitted.

Governments have begun developing institutional mechanisms that address parts of this problem. The United Kingdom's Algorithmic Transparency Recording Standard establishes a standardized way for public bodies to publish information about algorithmic tools, and government guidance defines algorithmic transparency as making information complete, open, understandable, accessible, and free. The standard became mandatory across central government in 2024. Canada requires departments within the scope of its Directive on Automated Decision-Making to assess automated systems through an Algorithmic Impact Assessment, publish relevant results, provide recourse, and meet requirements concerning transparency, quality, and accountability. These examples show that transparency is moving from voluntary ethical principle toward operational governance infrastructure.

Public algorithm registers provide another significant development. The Dutch government's Algorithm Register currently lists more than 1,500 algorithms and is designed to give visitors insight into impactful algorithms used by government organizations. Amsterdam's earlier municipal work helped shape the national register, while Dutch digital-government reporting emphasizes that publication alone is not sufficient and that algorithm registers should function as starting points for discussion about safeguards, risks, and responsibilities. Helsinki's AI Register similarly allows residents to view quick overviews or detailed information about municipal AI systems and to submit feedback, thereby adding an element of participation to transparency. These initiatives provide an important foundation, but they primarily operate as registers or structured records rather than fully developed decision-oriented transparency dashboards.



2. Objectives of the Study

2.1 To Define Citizen-Centered Algorithmic Transparency

The first objective is to distinguish citizen-centered transparency from conventional technical disclosure. The study examines what information an affected individual should reasonably be able to discover about a citizen-facing decision system, including its purpose, institutional owner, decision role, data categories, performance limitations, human-oversight arrangements, fairness assessments, and mechanisms for correction or appeal.

2.2 To Develop a Multidimensional Dashboard Evaluation Framework

The second objective is to create a framework for evaluating public algorithmic transparency dashboards. The framework covers disclosure completeness, citizen comprehensibility, actionability, performance transparency, fairness transparency, data provenance, governance accountability, accessibility, update integrity, and contestability.

2.3 To Compare Alternative Transparency Models

The third objective is to compare five representative transparency models: basic disclosure pages, technical algorithm registers, performance-monitoring dashboards, fairness-and-bias dashboards, and citizen-centered transparency dashboards.

3. Review of Literature

3.1 Algorithmic Transparency in Public Administration

Public-sector algorithmic transparency has developed from a broad accountability principle into a field involving disclosure standards, registers, public documentation, impact assessments, and system-specific explanation mechanisms. Veale's research with public-sector machine-learning practitioners showed that transparency operates within complex organizational contexts and that simply releasing information does not automatically resolve tensions involving trust, implementation, procurement, gaming, and professional practice. Veale, Van Kleek, and Binns subsequently identified a gap between technical fairness and transparency research and the organizational realities experienced by public-sector practitioners, emphasizing the need for usable tools capable of supporting risk identification, domain knowledge, and accountability in high-stakes environments.

This socio-technical understanding is important because transparency can fail through both excessive opacity and excessive complexity. A government can disclose hundreds of fields concerning model architecture, software dependencies, feature engineering, and mathematical metrics while still leaving ordinary citizens unable to determine why the system matters to them. Conversely, a highly simplified public description may become so generic that it provides no meaningful basis for scrutiny.

NIST's AI Risk Management Framework helps clarify the relationship between transparency and broader accountability. The AI RMF is a voluntary framework designed to integrate trustworthiness considerations into the design, development, use, and evaluation of AI systems. Its supporting materials emphasize governance, documentation, risk mapping, and measurement, while NIST specifically notes that transparency can reduce information asymmetry by providing visibility into AI pipelines, workflows, and organizational processes. A dashboard designed for public administration should

therefore make the institutional process surrounding the algorithm visible rather than focusing only on the technical model.

3.2 Algorithm Registers, Public Participation, and Actionability

Public algorithm registers have emerged as a widely discussed form of proactive algorithmic transparency. Helsinki's AI Register explicitly presents itself as a window into municipal AI systems and includes a feedback function that allows residents to participate in the city's development of human-centered AI. The Helsinki-Amsterdam white paper on public AI registers recommends transparency concerning purpose and impacts, accountability, datasets, data processing, non-discrimination, human oversight, risks, mitigation measures, and explainability.

The Dutch Algorithm Register provides a further example of transparency operating at national scale. As of August 2020, the public register contains more than 1,500 algorithm descriptions and focuses particularly on impactful algorithms and high-risk AI systems. Dutch government reporting in April 2020 stated that more than 1,300 algorithms had already been published by more than 500 government bodies and emphasized that registration is a starting point rather than an endpoint of responsible governance. This distinction is essential. Registers improve discoverability, but citizens may still require decision-specific information, current performance evidence, and access to redress.

Canada's model illustrates how transparency can be combined with structured risk assessment. Its Algorithmic Impact Assessment contains 65 risk questions and 41 mitigation questions and is mandatory for systems falling within the Directive on Automated Decision-Making. The Directive also requires publication of final AIA results in an accessible format and emphasizes recourse, quality assurance, and public reporting. The Canadian framework therefore demonstrates that public disclosure can be connected to governance controls rather than operating as an isolated communications exercise.

3.3 From Disclosure to Explainability and Contestability

Transparency is most consequential when an algorithm contributes to a decision that affects a particular individual. In these circumstances, general disclosure about the existence of a model may be less important than information that supports explanation and contestation.

A citizen-facing dashboard should therefore distinguish between system-level transparency and decision-level transparency. System-level transparency describes what the tool is, who owns it, what data it generally uses, how it was tested, and how it is governed. Decision-level transparency provides information about how the system contributed to a specific determination and what options exist for correction or reconsideration.

The distinction is consistent with NIST's description of transparency, explainability, and interpretability as related but separate properties. It is also reflected in Canada's Directive on Automated Decision-Making, which connects transparency and public reporting with recourse and procedural fairness rather than treating disclosure as a self-sufficient requirement.

The European Commission's July 2026 guidance on Article 50 of the AI Act further illustrates the growing importance of making people aware when particular forms of AI interaction or analysis occur. Article 50 transparency obligations became applicable on August 2, 2026 for specified categories of AI systems, although these requirements do not constitute a universal dashboard mandate for all public-



sector decision systems. Their significance for the present paper lies in the broader regulatory direction toward active notification rather than passive availability of information.

4. Research Methodology

4.1 Research Design

The study uses a conceptual-methodological design supported by simulation-based comparative analysis. No real citizen data, administrative case files, proprietary algorithms, or operational dashboard usage records were collected.

The evidence base combines public-sector algorithmic transparency scholarship with current government standards and responsible-AI frameworks. The primary governance references include the UK Algorithmic Transparency Recording Standard, Canada's Directive on Automated Decision-Making and Algorithmic Impact Assessment, the Dutch Algorithm Register, Helsinki's AI Register, NIST's AI Risk Management Framework, and current EU transparency guidance. The analysis models five transparency architectures and evaluates them through normalized 0–100 scores. The values are illustrative and are intended to demonstrate the structure of a dashboard evaluation method rather than measured public attitudes or observed system performance.

4.2 Dashboard Governance and Evaluation Dimensions

The proposed dashboard framework evaluates transparency through eight dimensions.

Table 1: Governance Framework for Citizen-Facing Algorithmic Transparency Dashboards

Dashboard Dimension	Core Question	Information to Be Disclosed	Governance Purpose
Purpose and Decision Role	Why is the system used and what influence does it have?	Administrative objective, decision stage, degree of automation	Prevents hidden or ambiguous algorithmic involvement
Data Provenance	What information enters the system?	Data categories, sources, update frequency, known quality limits	Supports scrutiny of relevance and accuracy
Performance	How reliably does the system operate?	Validation metrics, error rates, uncertainty, monitoring periods	Prevents unsupported claims of accuracy
Fairness and Bias	Do performance or outcomes differ across relevant groups?	Subgroup metrics, fairness assessments, identified limitations	Supports discrimination and equity scrutiny
Human Oversight	Where can human judgment intervene?	Reviewer role, override authority, escalation process	Clarifies whether automation is advisory or decisive



Dashboard Dimension	Core Question	Information to Be Disclosed	Governance Purpose
Accountability	Who owns and governs the system?	Responsible department, supplier role, contact authority	Prevents responsibility diffusion
Change and Incident History	Has the model changed or experienced material problems?	Versions, review dates, major incidents, corrective action	Converts disclosure into a living governance record
Explanation and Redress	What can an affected citizen do?	Explanation route, correction process, appeal mechanism, contact information	Converts transparency into practical actionability

Source: Author-developed framework informed by the UK ATRS, Canada's automated-decision governance requirements, Helsinki's AI Register, Dutch algorithm-register practice, and NIST AI governance principles.

4.3 Simulation Model

Five transparency models were constructed for methodological comparison.

- The Basic Disclosure Page provides a short description stating that an algorithm is used and identifies the responsible agency.
- The Technical Algorithm Register adds system design information, data sources, model description, and governance documentation but remains oriented primarily toward technically informed users.
- The Performance Monitoring Dashboard adds current operational statistics, error indicators, update dates, and performance trends.
- The Fairness and Bias Dashboard adds subgroup analysis, bias indicators, fairness assessments, and information regarding risk mitigation.
- The Citizen-Centered Transparency Dashboard combines these functions with plain-language explanations, decision-specific guidance, human-oversight information, institutional accountability, and accessible correction or appeal pathways.

5. Analysis

5.1 Comparative Evaluation of Dashboard Models

Table 2: Simulated Transparency Performance Across Dashboard Models

Dashboard Model	Disclosure Completeness	Citizen Actionability	Accountability Support	Simulated CTEI
Basic Disclosure Page	42	28	35	34.7
Technical Algorithm Register	68	35	55	51.6
Performance Monitoring Dashboard	76	58	66	66.3
Fairness and Bias Dashboard	84	67	78	75.7
Citizen-Centered Transparency Dashboard	92	91	94	92.1

Note: All values are simulated methodological examples. They do not represent survey results or evaluations of the UK, Dutch, Finnish, Canadian, or EU systems discussed in this paper.


The comparison demonstrates that transparency effectiveness is not determined by disclosure volume alone. The technical algorithm register achieves a disclosure score of 68 but a citizen-actionability score of only 35. This simulated difference represents a common governance problem: detailed technical records may be highly useful for auditors, journalists, researchers, and regulators while remaining difficult for affected citizens to translate into action.

The performance-monitoring dashboard improves actionability by showing whether the system operates as expected, but it still may not explain whether a particular citizen can challenge an individual decision. The fairness-and-bias dashboard produces stronger accountability by exposing subgroup performance and documented risks, although generalized fairness metrics still do not replace case-specific explanation.

The citizen-centered transparency dashboard achieves the highest simulated CTEI because it integrates system information with user-oriented explanation and redress. Its strongest feature is not simply that it publishes more fields, but that it connects information to the administrative relationship between the citizen and the institution.

5.2 Graphical Analysis of Transparency and Actionability

Graph 1: Simulated Transparency–Actionability Relationship Across Dashboard Models



Source: Author-generated simulation. Bubble size represents simulated accountability support.

The scatter graph demonstrates a positive relationship between disclosure completeness and citizen actionability, but the progression is not automatically proportional. The technical register displays substantially greater disclosure completeness than the basic disclosure page while producing only a modest improvement in actionability. This illustrates the difference between publishing information and enabling citizens to use it.

The citizen-centered transparency dashboard occupies the upper-right region because its information design is structured around citizen needs rather than system documentation alone. Its larger bubble represents stronger accountability support through responsible-authority identification, oversight information, audit history, and redress.

5.3 Proposed Citizen-Facing Dashboard Architecture

A strong transparency dashboard should begin with a plain-language system summary answering four immediate questions: What does this system do? Why does the government use it? Can it affect me? Does a human make the final decision? These questions should appear before technical documentation because they establish relevance.

The second layer should provide decision and data transparency. Citizens should be able to identify which categories of information the system uses, where those data originate, whether data are updated automatically, and which information can be corrected if inaccurate.

The third layer should disclose performance and limitation information. Accuracy should not be presented through a single percentage. Dashboards should distinguish false positive and false negative



risks where applicable, indicate evaluation dates, specify whether performance differs across relevant groups, and explain situations in which the system should not be relied upon.

The fourth layer should provide governance information, including responsible agency, system owner, supplier involvement, human-oversight arrangement, impact assessments, relevant audits, and current system version. UK and Dutch transparency initiatives demonstrate the value of standardized public information concerning system purpose and governance, while Canadian requirements illustrate how impact assessments can be published alongside automated decision governance.

The final layer should provide actionable citizen rights. Users should see how to obtain an explanation, correct inaccurate information, contact the responsible department, request human reconsideration, submit a complaint, or appeal a formal decision where applicable. Without this layer, transparency remains informational rather than remedial.

6. Discussion

6.1 Transparency Should Be Measured by Civic Usefulness Rather Than Information Volume

The analysis supports a central distinction between informational transparency and civic transparency. An organization may technically satisfy a disclosure requirement by publishing extensive documentation while providing little practical assistance to ordinary users. Government transparency should therefore be assessed according to whether the disclosed information enables citizens, journalists, civil-society groups, legislators, auditors, and regulators to understand how administrative power is being exercised. The UK ATRS is particularly instructive because its official definition emphasizes that algorithmic information should be complete, open, understandable, accessible, and free rather than merely published. The Dutch experience similarly shows that public registration is considered a starting point for debate about safeguards, risks, and institutional responsibility rather than an end in itself.

The simulated difference between the technical register and the citizen-centered dashboard illustrates this issue. A technically detailed register may provide high-quality documentation for specialist audiences while leaving a citizen unsure whether a specific algorithm affected a benefit application or what to do if an input variable is inaccurate. This should not be interpreted as an argument against technical documentation. Instead, the evidence should be layered for different audiences. Researchers and auditors require model-level information, regulators may require risk and compliance evidence, public officials need operational documentation, and affected citizens require clear explanations and redress pathways. A mature dashboard should serve these audiences through progressive disclosure rather than forcing all users into a single technical interface.

6.2 Performance and Fairness Transparency Must Be Dynamic

Static algorithm descriptions become incomplete when systems are updated, data distributions change, or operational performance shifts. A transparency dashboard should therefore function as a living governance interface rather than a publication artifact. Veale, Van Kleek, and Binns identified the importance of tracking changing data and system conditions in public-sector algorithmic governance. The Dutch Algorithm Register also treats government algorithm descriptions as ongoing records, while the current national platform includes recently modified entries and explicit register-maintenance infrastructure.



Dynamic transparency is especially important for performance and fairness. A model may initially achieve acceptable validation results and later perform differently because population characteristics, administrative procedures, input data, thresholds, or human usage patterns have changed. Dashboards should therefore display the date of the latest evaluation and avoid presenting old validation metrics as timeless characteristics of the system. Where meaningful, they should distinguish pre-deployment validation from operational monitoring.

6.3 Transparency Becomes Accountability Only When Citizens Can Act

The most important implication of the study is that transparency should connect directly to contestability. Citizens require more than awareness that an algorithm exists. They need to know whether the algorithm materially contributed to a decision, which information can be corrected, who can review the outcome, and which formal appeal or complaint procedure applies. Canada's Directive on Automated Decision-Making is notable because it connects transparency requirements to recourse and procedural fairness rather than treating public reporting as an isolated obligation. Helsinki's AI Register also embeds feedback into the transparency interface, demonstrating that public disclosure can be designed to support participation rather than one-way communication.

A citizen-centered dashboard should therefore treat redress as a core design element. The dashboard could provide a direct pathway from a system description to the responsible service unit, guidance on requesting human review, information on correcting source data, and the legal or administrative procedure governing formal appeals. These mechanisms should be integrated rather than placed in unrelated parts of a government website.

This action-oriented model also strengthens institutional learning. Complaints and appeals provide evidence about data errors, recurring misunderstandings, accessibility barriers, model limitations, and administrative practices. A dashboard can therefore become a two-way governance mechanism in which transparency helps citizens identify problems while citizen feedback helps institutions identify weaknesses in their systems.

7. Conclusion

Algorithmic transparency dashboards offer a practical opportunity to strengthen accountability as automated and AI-supported systems become embedded in citizen-facing public administration. Existing initiatives already demonstrate substantial progress. The United Kingdom has established a standardized algorithmic transparency regime and made the ATRS mandatory for central government, Canada connects automated decision-making with mandatory impact assessment and public reporting, Helsinki provides a public-facing AI register with feedback capabilities, and the Dutch national Algorithm Register now contains more than 1,500 published algorithm descriptions. These developments indicate that the institutional question is increasingly shifting from whether government algorithms should be disclosed toward how public information about them should be structured, maintained, and made useful. The present study argues that static disclosure alone is insufficient for citizen-facing decision systems.

Future research should empirically test algorithmic transparency dashboards with diverse citizen groups, including people with different levels of digital literacy, disability access needs, language backgrounds, and previous experience with public administration. Studies should examine whether



dashboards improve understanding, perceived procedural fairness, correction of administrative errors, successful appeals, public trust, and regulator effectiveness. Longitudinal research should also investigate whether published performance and fairness information remains current after deployment. The ultimate objective should not be transparency for its own sake. A well-designed dashboard should make algorithmically mediated public power more visible, understandable, reviewable, and contestable, thereby strengthening the institutional principle that technological complexity does not remove the government's obligation to explain and justify decisions made in the name of the public.

References

1. Diakopoulos, N. (2016). Accountability in algorithmic decision making. *Communications of the ACM*, 59(2), 56–62. <https://doi.org/10.1145/2844110>
2. 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.
3. Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press.
4. Burrell, J. (2016). How the machine “thinks”: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12. <https://doi.org/10.1177/2053951715622512>
5. Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973–989. <https://doi.org/10.1177/1461444816676645>
6. Selbst, A. D., & Barocas, S. (2018). The intuitive appeal of explainable machines. *Fordham Law Review*, 87(3), 1085–1139.
7. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
8. 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. <https://doi.org/10.1145/2939672.2939778>
9. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774.
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. <https://doi.org/10.1145/3236009>
11. Weller, A. (2019). Transparency: Motivations and challenges. In W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, & K.-R. Müller (Eds.), *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning* (pp. 23–40). Springer.
12. Floridi, L., Cowls, 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. *Minds and Machines*, 28, 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
13. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>



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. <https://doi.org/10.1080/01900692.2018.1498103>
15. Sun, T. Q., & Medaglia, R. (2019). Mapping the challenges of artificial intelligence in the public sector: Evidence from public healthcare. *Government Information Quarterly*, 36(2), 368–383. <https://doi.org/10.1016/j.giq.2018.09.008>
16. 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. <https://doi.org/10.1016/j.giq.2020.101493>
17. Valle-Cruz, D., Criado, J. I., Sandoval-Almazan, R., & Gil-Garcia, J. R. (2020). Assessing the public policy-cycle framework in the age of artificial intelligence: From policy to governance. *Government Information Quarterly*, 37(4), 101509. <https://doi.org/10.1016/j.giq.2020.101509>
18. Meijer, A. J., Curtin, D., & Hillebrandt, M. (2012). Open government: Connecting vision and voice. *International Review of Administrative Sciences*, 78(1), 10–29. <https://doi.org/10.1177/0020852311429533>
19. Grimmelikhuijsen, S. G., & Meijer, A. J. (2014). Effects of transparency on the perceived trustworthiness of a government organization: Evidence from an online experiment. *Journal of Public Administration Research and Theory*, 24(1), 137–157. <https://doi.org/10.1093/jopart/mus048>
20. OECD. (2019). *Recommendation of the Council on Artificial Intelligence*. OECD Publishing, Paris.