



Constitutional Sovereignty Doctrine: A Control-Based Architecture for Trustworthy AI, Data Sovereignty, and Accountable Digital Decision Systems

Celio Oliveira

Abstract: This research develops Constitutional Sovereignty Doctrine, or CSD, as a socio-technical governance framework for assessing whether a public institution retains sufficient lawful, operational, computational, evidentiary, and remedial control over an AI-enabled decision system for which it remains legally and constitutionally accountable. The doctrine extends beyond territorial data residency and infrastructure ownership by examining control across the decision chain, including lawful authority, data, infrastructure, models and inference, human judgement, evidentiary traceability, rights and review, and corrective capacity. The paper also introduces Accountability-Control Asymmetry (ACA), a diagnostic concept describing situations in which formal responsibility remains with a public institution while the practical capabilities required to supervise, explain, reconstruct, intervene in, or remedy a decision are distributed elsewhere. CSD is grounded in Constitutional Intelligence, Canadian administrative law, contemporary digital-sovereignty policy, research on meaningful human oversight and procedural fairness, and international AI-governance standards. The eight dimensions are expressed through a functional model, $CSD = f(A, D, I, M, H, E, R, C)$, and are treated as assessable institutional capabilities rather than as a predetermined numerical index. The research design therefore emphasises construct validity, expert assessment, architecture-based experiments, inter-rater reliability, and comparison with simpler sovereignty proxies such as data residency and infrastructure ownership.

Keywords: artificial intelligence, digital sovereignty, constitutional intelligence, public administration, accountable AI, human oversight, algorithmic governance, data governance

INTRODUCTION

Artificial intelligence is changing the architecture of public decision making before the language of public law and digital sovereignty has fully adapted to the distribution of technical control. Federal institutions remain legally responsible for administrative decisions even when important parts of the decision process depend on commercial cloud infrastructure, proprietary foundation models, external application programming interfaces, vendor-pushed updates, automated recommendation engines, and software components that departments do not operate directly. In that environment, the practical question is not limited to where data are stored. It is whether the institution that carries the legal and democratic responsibility for the decision still possesses the capabilities required to understand, supervise, interrupt, reconstruct, defend, and remedy it.

The point of departure for this research is Constitutional Intelligence: A Decision Architecture for Trustworthy AI Governance. That work treats public-sector AI as an institutional decision architecture rather than as a model considered in isolation. It links structured decision pathways, lawful authority, evidentiary inputs, audit trails, privacy

safeguards, vendor accountability, professional judgement, and rights review within the same governance system. The present research extends that architecture to a different question. When the computational and organisational components of a public decision are distributed across jurisdictions, suppliers, platforms, models, and human actors, how much effective constitutional control remains with the institution that is still accountable for the outcome?¹

This question has become more pressing because current digital sovereignty policy is itself moving away from a purely territorial model. The Government of Canada now defines digital sovereignty as the capacity to exercise autonomy over digital assets and services regardless of where technologies are developed, hosted, or supported. Its framework explicitly extends beyond storage and jurisdiction to operational resilience, system integrity, supply dependencies, and institutional control. Recent scholarship also characterises digital sovereignty as a socio-technical capacity involving legitimate authority, technological infrastructures, actors, and practices rather than as a simple property of physical location. The European Commission has moved in a similar direction through a cloud sovereignty framework that evaluates strategic, legal, data and AI, operational, supply chain, technological, security, compliance, and sustainability criteria.²³⁴

These developments support the premise of this research, but they do not resolve the public-decision problem. Digital-sovereignty frameworks generally assess technological autonomy, infrastructure, jurisdiction, supply chains, or data control at an organisational or national level. AI-governance frameworks, by contrast, provide more specific expectations concerning human oversight, traceability, fairness, risk, explanation, and accountability, but rarely frame these capabilities as conditions of sovereign institutional authority. Research on meaningful human oversight further demonstrates that formal human participation cannot by itself establish legitimate control where the decision-maker lacks the information, independence, competence, or practical authority required to exercise judgement⁵. The research gap therefore lies between sovereign capacity, institutional accountability, and meaningful decision control. Constitutional Sovereignty Doctrine is proposed as a bridge between these literatures by specifying the capabilities that an accountable public institution must retain, or be able reliably to exercise, across the lifecycle of an AI-enabled administrative decision.

1 Celio Oliveira, *Constitutional Intelligence: A Decision Architecture for Trustworthy AI Governance* (Altona, MB: FriesenPress, 2026), especially the DECIDE-PLUS methodology, FIVE-LAYER Architecture, IMPACT information architecture, and CHARTER RIGHTS DECISION TREE sections.

2 European Commission, "Sovereign Cloud Framework Explained," June 1, 2026, https://commission.europa.eu/news-and-media/news/sovereign-cloud-framework-explained-2026-06-01_en.

3 Blair Wang, Raffaele Fabio Ciriello, and Alexander Richter, "Digital Sovereignty," *Business & Information Systems Engineering* (2026), <https://doi.org/10.1007/s12599-026-01018-9>.

4 Government of Canada, "Digital Sovereignty: A Framework to Improve Digital Readiness of the Government of Canada," November 12, 2025, <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/cloud-services/digital-sovereignty/digital-sovereignty-framework-improve-digital-readiness.html>.

5 Daniel J. Escott, "An Intertraditionalist Approach to 'Humans in the Loop,'" in *Oxford Intersections: AI in Society*, ed. Philipp Hacker (Oxford University Press, 2025), <https://doi.org/10.1093/9780198945215.003.0208>.

The research deliberately avoids two claims that would be difficult to sustain. First, it does not claim that sovereignty is newly multidimensional. That argument already exists in contemporary policy and scholarship. Second, it does not claim that CSD has already been validated as a numerical index. The objective is narrower. This study defines the construct, derives observable dimensions from existing law, standards, governance practices, and prior research, and proposes an empirical design capable of validating, revising, or rejecting it. The framework is useful only if it can survive independent assessment and distinguish real differences in institutional control.

CONSTITUTIONAL INTELLIGENCE AS THE INSTITUTIONAL FOUNDATION

Constitutional Intelligence changes the unit of analysis from the algorithm to the decision system. The book develops four interlocking framework families: DECIDE-PLUS for structured decision logic, FIVE-LAYER Architecture for institutional accountability and stakeholder integration, IMPACT for executive knowledge transfer, and the CHARTER RIGHTS DECISION TREE for rights-sensitive AI governance. Although the frameworks were developed for trustworthy AI governance in Westminster public institutions such as Canada, their relevance to the present research is more general. They assume that technical performance is not sufficient to establish the legitimacy of a public decision and argue that public institutions must also be able to show the authority under which the system operates, the evidence relied upon, the path from evidence to recommendation, the role of human judgement, and the means by which a decision can be reviewed and corrected.⁶

A subsequent peer-reviewed synthesis further operationalised Constitutional Intelligence around rights-sensitive scoping, evidentiary traceability, stakeholder and rights assurance, stewardship and administrative continuity, and contestable oversight. That earlier article is the immediate scholarly antecedent to the present doctrine, but it did not formulate Constitutional Sovereignty Doctrine, Accountability-Control Asymmetry, or the Retained Accountability Principle developed here.⁷

For present purposes, the Constitutional Intelligence decision chain can be represented as follows: **Authority -> Data -> Inference -> Recommendation -> Human Judgement -> Decision -> Explanation -> Recourse**. This sequence is deliberately broader than a machine learning pipeline. An accurate prediction can still participate in a defective public decision if the institution: 1) lacks lawful authority for the use case, 2) cannot establish the provenance of the data, 3) cannot identify which model version was applied, 4) cannot show whether a human decision maker exercised independent judgement, or 5) cannot reconstruct the basis of the decision for review. The chain therefore links technical governance with administrative legitimacy.

This system's perspective has strong support in the accountability literature. Kroll and colleagues observed that accountability mechanisms developed for human decision makers often fail when applied directly to automated systems and argued for technical mechanisms capable of verifying compliance with legal and procedural standards. Raji and

⁶ Oliveira, Constitutional Intelligence.

⁷ Celio Oliveira, "Constitutional Intelligence as a Governance Architecture for Trustworthy Public-Sector AI," Open Access Journal of Data Science & Artificial Intelligence 4, no. 1 (2026): 000167, <https://doi.org/10.23880/oajda-16000167>.

colleagues later proposed an end-to-end internal auditing framework in which accountability is distributed across the development life cycle and documented at multiple stages rather than assessed only after deployment. These approaches are compatible with Constitutional Intelligence because each rejects the idea that accountability can be attached only to the final output. The process that creates the output must itself be observable and governable.⁸⁹

The same reasoning applies to data. The FAIR principles provide an important baseline for findability, accessibility, interoperability, and reusability. Constitutional Intelligence adds a public-law layer to those qualities. A department may have technically accessible and reusable data while still lacking adequate provenance, lawful authority, retention controls, lineage, versioning, or an evidentiary record sufficient to defend a consequential decision. CSD therefore separates data control from evidentiary control. The first concerns governance of the information entering and moving through the system. The second concerns whether the institution can later reconstruct the relevant decision state.¹⁰

Constitutional Intelligence also provides a reason to resist a purely efficiency-based view of AI governance. In a private system, a decision may be evaluated primarily in terms of performance, cost, user satisfaction, commercial risk, or profitability. A public institution faces additional constraints. Its exercise of authority must remain justifiable, reviewable, and consistent with the legal framework under which power was delegated. The Supreme Court of Canada has emphasised that administrative decision making is grounded in a culture of justification and that public power must be justified, intelligible, and transparent to those subject to it. This matters for AI because a technically high-performing system cannot substitute for an institution that has lost the ability to explain and justify the exercise of its own authority.¹¹

WHY TERRITORIAL LOCATION IS AN INCOMPLETE PROXY FOR CONTROL

Territorial location remains relevant to sovereignty, but it cannot carry the full analytical burden in distributed digital systems. The Government of Canada made this distinction explicit when it replaced an earlier public-cloud discussion focused largely on data residency and foreign jurisdiction with a broader digital sovereignty framework that includes operational resilience, system integrity, institutional control, workforce capacity, supply chains, and supplier dependence. The change is analytically important because it recognises that a system can be located in Canada and still be operationally dependent on actors or technologies outside the effective control of the institution using it.¹²

8 Inioluwa Deborah Raji et al., “Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing,” in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (New York: ACM, 2020), 33-44, <https://doi.org/10.1145/3351095.3372873>.

9 Joshua A. Kroll et al., “Accountable Algorithms,” *University of Pennsylvania Law Review* 165, no. 3 (2017): 633-705.

10 Mark D. Wilkinson et al., “The FAIR Guiding Principles for Scientific Data Management and Stewardship,” *Scientific Data* 3 (2016): 160018, <https://doi.org/10.1038/sdata.2016.18>.

11 Canada (Minister of Citizenship and Immigration) v. Vavilov, 2019 SCC 65, [2019] 4 S.C.R. 653, paras. 74-81, 95.

12 Government of Canada, “Digital Sovereignty.” See also Government of Canada, “Government of Canada White Paper: Data Sovereignty and Public Cloud,” 2018.

Cloud computing illustrates the problem clearly. Mitchell and Samlidis showed that government digital sovereignty concerns include offshore hosting, foreign cloud providers, public procurement, market concentration, and broader geopolitical dependencies. The European Commission has gone further in 2026 by operationalizing cloud sovereignty through measurable procurement criteria and assurance levels. Its framework includes legal and jurisdictional controls, operational autonomy, supply-chain visibility, technical documentation, open standards, and the capacity to maintain technology without continued dependence on a non-European supplier. The significance for this research is not that the European model should be imported into Canadian constitutional analysis. It is that contemporary public-sector sovereignty practice increasingly treats control as measurable through capabilities and dependencies, not location alone.¹³¹⁴

AI creates an even more difficult form of separation because decision-relevant control can migrate while the underlying data remain in the same jurisdiction. A department can retain Canadian data residency yet rely on a proprietary model whose updates are controlled by a supplier, whose evaluation data are unavailable, whose inference environment is partly opaque, or whose output becomes functionally determinative in a high-volume process. In such a case, the important sovereignty loss may occur at the level of model governance, inference, evidentiary reconstruction, or decision dependence rather than at the level of data storage.

This leads to four distinctions that the empirical study should test rather than merely assert. Location is evidence about control, but location is not identical to control. Ownership is evidence about control, but ownership is not identical to control. Contractual responsibility is evidence about accountability, but it does not guarantee operational capability. Human presence in a workflow is evidence of possible oversight, but it does not establish meaningful human judgement. Each proxy becomes useful only when connected to the underlying capabilities that the institution actually possesses.

The central research implication is that sovereignty should be examined at the point where authority is exercised. For public-sector AI, that point is the decision architecture. The institution may have strong control at one node and weak control at another. It may control data but not the model, control infrastructure but not evidence, retain formal decision authority but lack meaningful capacity to depart from automated recommendations, or remain legally accountable while contractual and technical dependencies make intervention difficult. CSD is designed to make those differences visible.

ACCOUNTABILITY-CONTROL ASYMMETRY

The first original construct proposed here is Accountability-Control Asymmetry. It describes a condition in which the institution formally responsible for an AI-mediated decision lacks one or more material capabilities required to inspect, intervene in, reconstruct, explain, or

¹³ European Commission, Cloud Sovereignty Framework: Implementation Guidance (Brussels: European Commission, 2026); European Commission, "Commission Advances Cloud Sovereignty through Strategic Procurement," April 17, 2026.

¹⁴ Andrew D. Mitchell and Theodore Samlidis, "Cloud Services and Government Digital Sovereignty in Australia and Beyond," *International Journal of Law and Information Technology* 29, no. 4 (2021): 364-394, <https://doi.org/10.1093/ijlit/eaac003>.

remedy that decision. This is not identical to the general accountability gap discussed in the algorithmic audit literature. The accountability gap concerns the difficulty of making AI systems answerable.

Accountability-Control Asymmetry focuses more narrowly on a separation between retained responsibility and migrated capability. The relationship is expressed conceptually as follows:

$$ACA = A_f - C_e$$

In this expression, A_f represents formal accountability obligations and C_e represents effective control capability. The expression is not presented as a validated subtraction producing a scalar value. It is a diagnostic model. Formal accountability can be identified through statute, delegated authority, administrative responsibility, policy, contract, or review obligations.

Effective control can be observed through access rights, model governance, auditability, the capacity to suspend or override the system, portability, termination rights, evidence retention, and remedial authority.

The legal rationale for examining this asymmetry predates contemporary AI. Citron argued that automated systems can weaken due process where code becomes embedded in adjudication without equivalent mechanisms for notice, participation, accountability, and correction. Canadian administrative law reaches the issue from a different doctrinal tradition, but with a related concern for fair procedure. *Baker v. Canada* established that the content of procedural fairness is contextual and depends on factors including the nature of the decision, the statutory scheme, the importance of the decision to the affected person, legitimate expectations, and procedural choices. Vavilov later reinforced the relationship between justification, transparency, intelligibility, and administrative legitimacy.¹⁵¹⁶¹⁷

These cases do not establish a legal rule called Accountability-Control Asymmetry, and this research does not present them as doing so. Their importance is structural. They show why a public institution must retain enough knowledge and control to discharge the obligations attached to its authority. A department cannot meaningfully provide reasons, respond to review, or correct a defective decision if the operational evidence needed to do so is unavailable or controlled entirely by another actor. In that sense, technical control can become a precondition for legal accountability without being identical to legal accountability.

The Government of Canada Directive on Automated Decision-Making provides a contemporary operational bridge. The Directive requires institutions within scope to assess automated decision systems, apply impact-based requirements, support explanation and recourse, undertake appropriate review, and use automation in a manner compatible with transparency, accountability, legality, and procedural fairness. The Algorithmic Impact Assessment asks questions across business processes, data, system designs, algorithms,

15 Vavilov, 2019 SCC 65, paras. 74-81, 95, 127, 133.

16 *Baker v. Canada (Minister of Citizenship and Immigration)*, [1999] 2 S.C.R. 817, paras. 21-28.

17 Danielle Keats Citron, "Technological Due Process," *Washington University Law Review* 85, no. 6 (2008): 1249-1313.

decisions, impact, and mitigation. This is significant because federal practice already treats accountable automation as a multidimensional institutional problem rather than a model-performance problem alone.¹⁸¹⁹

CONSTITUTIONAL SOVEREIGNTY DOCTRINE

Constitutional Sovereignty Doctrine (CSD) holds that where a public institution retains legal and constitutional responsibility for a consequential digital decision, it must retain, or be able reliably to exercise, the practical capabilities necessary to govern that decision. Those capabilities include lawful authority, control over relevant data and infrastructure, sufficient authority over models and inference, meaningful human decision authority, evidentiary traceability, protection of rights and review, and the capacity to intervene, correct, enforce, and maintain continuity.

The doctrine does not require technological self-sufficiency or exclusive institutional ownership of every component. Public institutions can rely on external infrastructure, commercial suppliers, shared platforms, federated data, and proprietary models. The relevant question is whether those dependencies preserve or materially impair the institution's continuing capacity to exercise the authority and discharge the obligations for which it remains accountable. The functional representation is:

$$\text{CSD} = f(\text{A}, \text{D}, \text{I}, \text{M}, \text{H}, \text{E}, \text{R}, \text{C})$$

where **A** represents lawful authority; **D**, data control; **I**, infrastructure and operational control; **M**, model and inference control; **H**, human and decision authority; **E**, evidentiary traceability; **R**, rights, transparency, and review; and **C**, corrective, enforcement, and continuity capability.

The formula is not presented as a validated numerical equation or as evidence that the dimensions carry equal weights. It is an analytical representation of the domains through which constitutional control can be examined. Empirical validation should determine whether these dimensions are observable, sufficiently distinct, practically useful, and capable of differentiating institutional architectures with materially different levels or forms of accountable control. This formulation distinguishes doctrine from measurement. CSD establishes the underlying proposition and identifies the relevant control capabilities. Empirical methods may subsequently operationalise those capabilities through indicators, thresholds, profiles, or other assessment methods, but the validity of the doctrine does not depend on reducing constitutional sovereignty to a universal percentage or scalar score.

THE EIGHT DIMENSIONS OF CSD

The first dimension, lawful and institutional authority, asks whether the institution has a defensible legal and administrative basis for the use of AI in the particular decision context. This is not a general question about whether the organisation is allowed to use technology.

18 Treasury Board of Canada Secretariat, "Algorithmic Impact Assessment," Open Government Portal, accessed October 2, 2026, <https://open.canada.ca/aia-eia-js/?lang=en>.

19 Treasury Board of Canada Secretariat, "Directive on Automated Decision-Making," modified June 24, 2025, <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592>.

It concerns whether the specific use of data, inference, automation, and delegated discretion can be traced to lawful institutional authority. Indicators may include the existence of a statutory or policy basis, the identity of the accountable decision maker, the scope of delegated authority, the documented purpose of the system, and the relationship between the technology and the decision that the institution is authorised to make.

The second dimension, data control, examines the institution's capacity to govern the information entering and moving through the system. Relevant indicators include provenance, lawful collection and access, data quality, lineage, retention, deletion, movement across environments, access control, and the ability to identify which data were used for a particular decision. The aim is not to equate sovereignty with exclusive possession of data. Interoperable and federated systems may preserve substantial control if access, use, transformation, and movement remain governable and auditable.

The third dimension, infrastructure and operational control, concerns who can operate, configure, secure, maintain, recover, and materially alter the environment in which the system runs. Contemporary sovereignty frameworks support treating these questions independently of geography. Canada's framework emphasises system continuity, supplier dependence, assurance, and the ability to maintain services in a globally connected environment. The European Commission's cloud sovereignty criteria similarly examine operational autonomy, technical documentation, supply chains, open standards, and the capacity to maintain services without continuous external dependence.²⁰²¹

The fourth dimension, model and inference control, is especially important for AI because model governance can be separated from both data control and infrastructure control. Indicators can include whether the institution determines which model version is deployed, can prevent unapproved updates, can define or constrain inference parameters, can independently evaluate performance, can stop inference, can replace the model, can inspect relevant documentation, and can detect material changes in behaviour. A public institution may own none of the model weights and still retain meaningful control if enforceable governance mechanisms provide equivalent practical capability. Conversely, nominal ownership may be weak evidence of control if the institution lacks the expertise or technical access required to exercise it.

The fifth dimension, human and decision authority, asks who actually determines the consequential outcome. It distinguishes legal attribution from practical influence. A person who formally signs a decision may have little real authority if an automated score is treated as presumptively correct, if contradictory evidence is difficult to enter, if workload makes independent review unrealistic, or if the person cannot see the information needed to challenge the recommendation. The relevant measure is therefore not whether a human appears in the workflow. It is whether the human can exercise informed, independent, and consequential judgement. This distinction is also consistent with legal scholarship questioning whether the mere presence of a human in an AI-mediated process constitutes meaningful oversight. In adjudicative settings, meaningful human involvement has been linked to deeper requirements of procedural fairness, judicial independence, the right to be heard, and the duty to give reasons. The relevant governance question is therefore not

20 European Commission, Cloud Sovereignty Framework: Implementation Guidance.

21 Government of Canada, "Digital Sovereignty."

simply whether a human remains “in the loop”, but whether the institutional and technological architecture preserves the conditions under which independent human judgement can actually be exercised.

The sixth dimension, evidentiary traceability, concerns the institution’s capacity to reconstruct the relevant decision state. This includes the data relied upon, model or rules applied, software version, time of execution, relevant parameters, recommendation produced, human interventions, final decision, and subsequent corrections. Raji and colleagues’ lifecycle audit model provides strong support for documentation across stages of system development and deployment because post-deployment harms can otherwise become difficult to trace to their source. The EU AI Act similarly requires automatic logging for high-risk systems, reflecting the importance of records for monitoring and accountability.²²²³

The seventh dimension, rights, transparency, and review, examines whether people affected by a system can understand the role of automation sufficiently to exercise their rights and obtain meaningful review. Selbst and Powles argue that information about automated decision logic should be interpreted functionally, so that disclosure enables the exercise of substantive rights rather than satisfying a purely formal transparency requirement. The Council of Europe Framework Convention likewise links documentation and transparency to procedural safeguards, complaints, challenges, and remedies. Internal observability is therefore not enough. A system may be highly transparent to engineers while remaining practically unchallengeable to a citizen.²⁴²⁵

The eighth dimension, corrective, enforcement, and continuity capability, asks whether the accountable institution can act when something goes wrong. This includes the ability to suspend the system, override or reverse an output, correct data, change a model or rule, investigate an incident, provide a remedy, move to another supplier, and continue essential services when a dependency fails. The OECD AI Principles explicitly call for mechanisms to override, repair, or safely decommission AI systems when they present undue harm or undesired behaviour and also connect accountability to lifecycle traceability. This dimension turns accountability from a descriptive label into an institutional capability.²⁶

HUMAN JUDGEMENT AND THE LIMITS OF NOMINAL OVERSIGHT

Human oversight requires separate treatment because it is frequently invoked as if the presence of a person were sufficient to solve the accountability problem. Constitutional Intelligence does not support that assumption. Meaningful human judgement depends on authority, information, competence, time, and the practical ability to disagree with the system. A workflow in which a public servant can technically reject a recommendation but

22 European Union, Regulation (EU) 2024/1689, art. 12.

23 Raji et al., “Closing the AI Accountability Gap,” 33-44.

24 Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (Strasbourg: Council of Europe, 2024).

25 Andrew D. Selbst and Julia Powles, “Meaningful Information and the Right to Explanation,” *International Data Privacy Law* 7, no. 4 (2017): 233-242, <https://doi.org/10.1093/idpl/ix022>.

26 OECD, “OECD AI Principles,” accessed October 2, 2026, <https://www.oecd.org/en/topics/ai-principles.html>.

lacks access to contradictory evidence, does not know which model version generated the result, or faces procedural pressure to follow the automated output should not be treated as equivalent to independent human decision making.

The empirical literature supports a careful rather than categorical claim. Alon-Barkat and Busuioc conducted three experimental studies, including a large sample of civil servants, and did not find a general pattern of automatic adherence to algorithmic advice. They did find evidence in one study of selective adherence when advice aligned with group stereotypes and cautioned that repeated exposure to high-performing systems may alter patterns of deference over time. Their findings are useful precisely because they prevent an exaggerated proposition. The risk is not that AI necessarily eliminates human discretion. The issue is whether institutional conditions preserve the quality and independence of that discretion.²⁷

Other experimental work shows why those conditions matter. Cecil and colleagues conducted five preregistered experiments involving 1,403 participants and found that incorrect AI advice reduced performance because participants often failed to dismiss it. Their explainability interventions did not reliably eliminate over-reliance. This suggests that explanation alone cannot be treated as a substitute for decision authority, evidence access, or review capacity. In CSD, human oversight should therefore be operationalised through observable powers: access to the underlying evidence, ability to obtain alternative information, authority to override, freedom to depart from the recommendation, documented reasons for intervention, and access to escalation or specialist review.²⁸

The EU AI Act takes a compatible approach for high-risk systems. Article 14 requires systems to be designed so that natural persons can effectively oversee them during use and links oversight to the prevention or minimisation of risks to health, safety, and fundamental rights. The significance for CSD is that oversight is defined functionally. Effective oversight depends on system design and the capacities available to the person, not on a formal statement that a human remains responsible.²⁹

TRACEABILITY, REASONS, AND EVIDENTIARY CONTROL

Traceability is the point at which technical observability becomes administrative evidence. If an institution cannot reconstruct a decision, it cannot reliably investigate an error, show that relevant information was considered, respond to an affected person, support internal review, or defend the exercise of public authority. This is why CSD treats evidentiary control as more than logging. Logs are useful only if they preserve the information needed to understand what occurred and if the institution can access and interpret them when accountability is engaged.

Canadian administrative law provides a strong conceptual basis for this emphasis. Vavilov explains that reasons help demonstrate that a decision was made in a fair and lawful

27 Saar Alon-Barkat and Madalina Busuioc, "Human-AI Interactions in Public Sector Decision Making: 'Automation Bias' and 'Selective Adherence' to Algorithmic Advice," *Journal of Public Administration Research and Theory* 33, no. 1 (2023): 153-169, <https://doi.org/10.1093/jopart/muac007>.

28 Julia Cecil et al., "Explainability Does Not Mitigate the Negative Impact of Incorrect AI Advice in a Personnel Selection Task," *Scientific Reports* 14 (2024): 9736, <https://doi.org/10.1038/s41598-024-60220-5>.

29 European Union, Regulation (EU) 2024/1689, art. 14.

manner and that administrative legitimacy is connected to justification, transparency, and intelligibility. CSD does not transform that jurisprudence into a technical specification. It draws the narrower inference that a public institution cannot reliably satisfy those expectations if the system architecture prevents it from reconstructing the decision path. In high-impact contexts, evidentiary traceability becomes part of the infrastructure that makes justification possible.³⁰

There is also a temporal dimension. AI systems change. Models are retrained, prompts and retrieval sources are altered, data distributions drift, vendors update services, and policies change. A decision that can be reproduced today may not be reproducible six months later if the relevant version state has not been preserved. This creates a specific research requirement: CSD indicators should measure not only current observability but also whether the institution preserves enough historical state to reconstruct past decisions after the technology has changed.

This is one reason the proposed framework should be applied to the entire decision chain rather than only to the deployed model. A complete evidentiary record may require data lineage, model versioning, system configuration, decision rules, timestamps, human interventions, and final outcomes. It may also require documentation of vendor changes and the conditions under which the institution accepted those changes. The object of evidence is therefore the decision architecture, not simply the algorithm. Evidentiary control also requires a distinction between traceability and verifiability. Traceability asks whether the institution can reconstruct how an output was produced. Verifiability asks whether the responsible professional or institution can independently establish that material relied upon in reaching the decision is accurate, authentic, and adequately supported. This distinction is particularly important for generative AI, where an output may be traceable to a system interaction while still containing unsupported factual or legal material. In consequential decision environments, retained accountability therefore requires not only access to an output or explanation, but a practical capacity to verify the evidence on which institutional action depends.

CONTESTABILITY AND REMEDY AS SYSTEM PROPERTIES

Contestability is often described as a legal or procedural right, but it has technical prerequisites. A person cannot effectively challenge a decision if the institution cannot identify which information mattered, which system was used, how the recommendation influenced the outcome, or who had authority to change the result. The Council of Europe Framework Convention requires documentation and information sufficient to support challenges, complaints, procedural safeguards, and remedies in connection with AI systems. This demonstrates the connection between rights and system design.³¹

The same distinction appears in the explanation literature. Selbst and Powles argue for a functional understanding of meaningful information that enables an affected person to exercise rights rather than merely receive formal technical disclosure. For CSD, the practical implication is that explanation quality should be evaluated against the decision

³⁰ Vavilov, 2019 SCC 65, paras. 77-81, 95-96.

³¹ Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law.

context. A technical description of a model may be irrelevant to a person seeking to know why a benefit was denied. Conversely, a simplified explanation that omits the decisive factor may be easy to read but inadequate for review. The institution must retain enough control to produce the kind of explanation the legal and administrative context requires.³²

Remedy extends the analysis beyond explanation. An institution may be able to explain a wrong decision while still lacking the technical or contractual power to correct the underlying problem. Corrective capacity therefore includes both case-level remedy and system-level intervention. Case-level remedy may involve reconsideration, correction of data, or reversal of an outcome. System-level intervention may involve suspending the model, changing rules, replacing a supplier, rolling back a version, or preventing recurrence. This distinction is central to the Retained Accountability Principle developed below.

VENDOR DEPENDENCE, INTEROPERABILITY, AND EXIT CAPACITY

Public institutions increasingly exercise authority through technology they do not build themselves. Vendor participation does not by itself reduce sovereignty. The relevant question is whether the institutional relationship preserves control. A commercial platform may provide stronger security, resilience, and operational capability than an internally operated system. At the same time, concentration, proprietary interfaces, inaccessible documentation, restricted audit rights, or difficult exit conditions can create dependencies that weaken the institution's ability to govern the public function over time.

Current sovereignty policy supports treating this as a measurable problem. The Government of Canada identifies concentration risk, global supply-chain exposure, service continuity, supplier diversification, and transparency about operational control as sovereignty considerations. The European Commission's 2026 cloud framework similarly includes supply-chain visibility, audit rights, technical documentation, non-proprietary interfaces, open standards, and the ability to operate and sustain services without continued reliance on non-European vendors. These policy developments justify including vendor and exit capabilities within the operational and corrective dimensions of CSD.³³³⁴

Interoperability is particularly important because it can preserve sovereignty without requiring isolation. Constitutional Intelligence has consistently favoured access over unnecessary copying and treats interoperability as a governance capability rather than a synonym for unrestricted data movement. In CSD, interoperability should be tested through portability of data, exportability of configurations and logs, use of documented interfaces, ability to substitute components, and preservation of evidentiary records during transition. A system can be internationally distributed while retaining substantial institutional control if those capabilities are enforceable and technically real.

Earlier Canadian work on public-sector data interoperability likewise treated interoperable public data as strategic infrastructure for responsible AI adoption, digital sovereignty, and coherent action across federal, provincial, and territorial jurisdictions. That analysis supports treating interoperability and coordinated data governance as

³² Selbst and Powles, "Meaningful Information and the Right to Explanation," 233-242.

³³ European Commission, Cloud Sovereignty Framework: Implementation Guidance.

³⁴ Government of Canada, "Digital Sovereignty."

enabling conditions for sovereign institutional capacity. The present doctrine narrows the inquiry further by asking which control capabilities must remain practically exercisable at the level of a consequential public decision.³⁵

This point also limits a simplistic localisation approach. The European Commission has stated that even non-European technologies can meet minimum sovereignty requirements when operated within an appropriate framework. That observation is directly relevant to the Canadian context, where public institutions will continue to depend on global technology markets. The research question is therefore not whether external technology should be excluded. It is what conditions allow the institution to use external technology without losing the practical capacity to discharge public responsibilities.³⁶

THE RETAINED ACCOUNTABILITY PRINCIPLE

The second normative contribution of this research is the Retained Accountability Principle: where a public institution retains legal responsibility for a consequential decision, delegation of computational functions should not extinguish its practical capacity to inspect, contest, intervene in, reconstruct, and remedy the resulting decision. The principle is proposed here as a governance proposition. It is not presented as an existing rule of Canadian law or international law.

The principle is nevertheless grounded in convergent requirements. The Government of Canada expects automated decision systems to operate consistently with accountability, legality, transparency, and procedural fairness. OECD principles connect accountability to lifecycle traceability and to the ability to override, repair, or decommission systems. The Council of Europe connects AI governance to oversight, procedural safeguards, complaints, and remedies. NIST treats accountability and transparency as cross-cutting socio-technical characteristics. Taken together, these sources justify examining retained capability as a governance requirement even though none formulates the principle in these terms.³⁷³⁸³⁹⁴⁰

The principle is intentionally stronger than a requirement to assign responsibility.

Responsibility can be allocated on paper while the responsible institution lacks operational access or leverage. Retained accountability instead asks whether the actor bearing responsibility can still do the things that responsibility requires. This is where the principle intersects with Accountability-Control Asymmetry. The asymmetry diagnoses the gap. The Retained Accountability Principle provides the normative direction for reducing it.

35 Alain Dudoit, Tony Labilloy, and Celio Oliveira, Public Sector Data Interoperability and Responsible Adoption of AI in Canada: Strategic Overview and Call to Action, CIRANO 2026PR-02 (2026), <https://doi.org/10.54932/PUUN9151>.

36 European Commission, "Commission Advances Cloud Sovereignty through Strategic Procurement," April 17, 2026.

37 National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework, 12.

38 Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law.

39 OECD, "OECD AI Principles."

40 Treasury Board of Canada Secretariat, "Directive on Automated Decision-Making."

EMPIRICAL OPERATIONALISATION OF CSD

CSD should be treated as a doctrinal and diagnostic framework with empirically testable dimensions, not as a completed index. The empirical design should therefore begin by testing whether the eight proposed domains are valid, observable, sufficiently distinct, and useful for differentiating institutional architectures. Design science research is appropriate because the study develops a structured assessment artefact derived from the doctrine and then evaluates whether that artefact performs its intended analytical function. A mixed-methods design is preferable because lawful authority, technical control, organisational capability, human judgement, evidence, rights, and remedy are unlikely to be captured adequately through one method.

The first phase is systematic construct derivation. Each dimension should be mapped to authoritative sources and translated into observable indicators. For lawful authority, indicators can examine whether the system has a documented legal and policy basis, a clearly identified accountable decision maker, and defined limits on delegated automation. For data control, indicators can examine provenance, access, lineage, movement, retention, deletion, and versioning. For infrastructure control, they can examine administrative access, continuity, recovery, configuration authority, and dependency on external operators.

For model and inference control, indicators can test whether the institution knows which model version is in use, can prevent unapproved updates, can evaluate behaviour independently, can adjust relevant parameters, can stop inference, can replace the model, and can obtain sufficient documentation to assess changes. For human and decision authority, indicators can test whether a person has access to relevant evidence, can reject or modify a recommendation, can seek an alternative analysis, and can do so without disproportionate procedural barriers. For evidentiary control, indicators can test whether the institution can reproduce the decision record at a later date.

Rights and review indicators should examine notice, explanation, channels for challenge, access to a human review process, and the ability of the reviewer to change the outcome.

Corrective and continuity indicators should examine suspension, rollback, incident response, portability, supplier exit, remedy, and service continuity. These indicators should be framed as verifiable capabilities rather than aspirational statements. A statement that a vendor is accountable is weaker evidence than a contractual and technical arrangement that gives the institution access to logs, documented change controls, and an enforceable exit path.

EXPERT VALIDATION AND ARCHITECTURE EXPERIMENT

The second phase should use expert content validation. An interdisciplinary panel of approximately 20 to 40 participants could include AI engineers, cybersecurity specialists, public administrators, data governance practitioners, privacy professionals, administrative lawyers, procurement specialists, and AI assurance professionals. A Content Validity Ratio or Content Validity Index can be used to identify indicators that experts consider essential or relevant. A two-round Delphi process can then test whether disagreement reflects

ambiguity in the indicator, disciplinary differences, or a substantive challenge to the construct.

The third phase should use controlled architecture scenarios. The scenarios should vary one or more control conditions while holding other features constant. A fully institution-controlled architecture can serve as a reference condition. A second architecture can keep data and infrastructure domestically while relying on a proprietary model with supplier-controlled updates and incomplete logs. A third can use an internationally distributed commercial ecosystem with weak audit rights and nominal human approval. A fourth can be federated across jurisdictions while preserving institutionally controlled encryption, strong access controls, auditable models, independent logs, meaningful human decision authority, portability, and enforceable recourse.

The purpose of the scenarios is not to produce a predetermined ranking. They are designed to test whether the dimensions behave as expected. If experts consistently assess the federated but strongly governed architecture as retaining more effective institutional control than the domestically hosted but externally controlled architecture, territorial localisation would be shown to be an incomplete proxy for the target construct. If the result is reversed, or if experts cannot distinguish the architectures reliably, the framework would require revision. Falsifiability is a strength of the research, not a threat to it.

Three hypotheses follow from the conceptual model. H1: territorial data location alone will explain less variation in expert assessments of institutional digital control than a multidimensional model incorporating operational, technical, evidentiary, and decision authority. H2: architectures retaining institutional control over model configuration, audit evidence, decision authority, and remediation will receive higher assessments of effective sovereignty than architectures possessing equivalent territorial data localisation without those capabilities. H3: nominal human participation will not significantly increase assessments of accountable control where the human actor lacks sufficient evidence, meaningful override authority, or the practical ability to depart from AI recommendations.

Inter-rater reliability can be tested through an intraclass correlation coefficient where the study design supports it. Internal consistency can be explored through McDonald's omega or Cronbach's alpha only if the final indicators are theoretically expected to reflect a common latent construct. Factor analysis should be reserved for a sample large enough to justify it. Sensitivity analysis can test whether removing a particular control changes expert assessments in the expected direction. Most importantly, CSD should be compared against simpler baselines such as data residency, infrastructure ownership, legal jurisdiction, or nominal human oversight. The new framework should only be retained if the additional complexity produces meaningful explanatory or discriminatory value.

AVOIDING FALSE PRECISION AND COMPENSATORY SCORING

The first empirical study should not reduce CSD to a percentage. A result such as 73.6 per cent would imply a precision that the construct has not yet earned. It could also conceal critical deficiencies. A system with excellent domestic infrastructure and data governance but no lawful authority should not receive a respectable overall result simply because high scores in other dimensions compensate mathematically for the defect. The same problem

arises where a system has strong technical controls but no meaningful remedy or where a decision is formally human but practically determined by an unreviewable automated score.

A profile-plus-threshold model is therefore more defensible for initial validation. Lawful authority, meaningful decision control, evidentiary traceability, and intervention or remedy should be examined as candidate critical conditions. Other dimensions can remain gradual. The empirical work should test whether experts in fact treat those dimensions as non-compensatory. NIST's observation that AI trustworthiness depends on interacting characteristics and may be constrained by weak characteristics provides an external rationale for this cautious approach.⁴¹

This restraint is deliberate. CSD should support quantitative measurement only where the evidence justifies it. Its primary function is doctrinal and diagnostic: to identify where authority and control reside, where capabilities have migrated, which dependencies matter, and whether the accountable institution still has the practical capacity to act. Empirical validation may strengthen, refine, or challenge the proposed dimensions without requiring constitutional sovereignty to be collapsed into a single numerical index.

RESEARCH ORIGINALITY AND RELATIONSHIP TO EXISTING WORK

The originality claim should be deliberately narrow. The paper should not state that digital sovereignty has never been conceptualised as multidimensional, nor that current frameworks are confined to data residency. Wang, Ciriello, and Richter explicitly frame digital sovereignty through interdependent actors, technological infrastructures, and socio-technical practices.

Canada's framework includes legal, security, privacy, workforce, supply-chain, operational, and institutional factors. The European Commission now measures cloud sovereignty across eight objective categories and assurance levels. These developments are prior art for the general proposition that digital sovereignty is multidimensional.⁴²

The new contribution is the decision-centred integration of that sovereignty literature with public-sector AI accountability. CSD moves the unit of analysis from national technological capacity or cloud infrastructure to the chain through which a public decision is authorised, supplied with evidence, computed, interpreted, approved, explained, challenged, and remedied. It then proposes a method for testing whether the responsible institution retains effective control at each stage. CSD therefore extends, rather than restates, two prior strands of the author's published work: the decision-system governance architecture developed through Constitutional Intelligence and the interoperability analysis linking responsible AI adoption, digital sovereignty, and coordinated public data governance. The original contribution of this article is the retained-control doctrine itself, its eight operational dimensions, Accountability-Control Asymmetry, the Retained Accountability Principle, and the proposed empirical validation architecture.⁴³

41 National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework, 12.

42 Wang, Ciriello, and Richter, "Digital Sovereignty."

43 See Oliveira, "Constitutional Intelligence as a Governance Architecture for Trustworthy Public-Sector AI"; Dudoit, Labilloy, and Oliveira, Public Sector Data Interoperability and Responsible Adoption of AI in Canada.

The second new contribution is Accountability-Control Asymmetry, which isolates the condition in which legal or administrative responsibility is retained while capabilities needed to discharge that responsibility are located elsewhere. The third is the Retained Accountability Principle, which proposes that delegation should not extinguish the institution's practical capacity to inspect, contest, intervene in, reconstruct, and remedy the decision. The fourth is empirical. The research does not simply propose a richer sovereignty framework. It requires the framework to demonstrate added value against simpler proxies and permits the evidence to show that some proposed dimensions are weak, redundant, or context-specific.

This separates the present research from Constitutional Intelligence itself. Constitutional Intelligence provides a governance architecture for trustworthy AI and establishes the importance of structured decisions, rights protection, evidence, human judgement, institutional accountability, vendor evaluation, and continuous governance. CSD addresses the next problem: when that architecture is distributed across jurisdictions, suppliers, models, platforms, and human actors, what effective constitutional control remains with the institution that is still accountable for the outcome?

HUMAN OVERSIGHT, VERIFICATION, AND PROCEDURAL JUSTICE

A control-based account of constitutional sovereignty must distinguish formal responsibility from the practical conditions required to exercise authority legitimately. The distinction becomes especially important in adjudicative and other consequential decisions. A human decision maker may remain formally responsible yet exercise little meaningful control if relevant evidence is unavailable, machine-generated material cannot be independently assessed, departure from an automated recommendation is constrained, or institutional conditions make independent judgement impracticable.

Research on meaningful human oversight links human participation to procedural fairness, independence, the right to be heard, and reasoned decision making.⁴⁴ The question is not simply whether a person remains 'in the loop'. It is whether the surrounding institutional and technological architecture preserves the information, competence, discretion, authority, and procedural space needed for independent judgement. Verification poses a related problem.

Receiving an AI-generated output does not discharge professional or institutional responsibility where the factual, legal, or evidentiary material on which it relies cannot be independently checked.⁴⁵ Access-to-justice research exposes a further limit. Strong institutional control does not, by itself, make a system accessible, procedurally fair, non-discriminatory, or capable of meaningful review.⁴⁶ Courts, tribunals, and administrative

44 Daniel J. Escott, "An Intertraditionalist Approach to 'Humans in the Loop,'" in *Oxford Intersections: AI in Society*, ed. Philipp Hacker (Oxford University Press, published 23 October 2025), <https://doi.org/10.1093/9780198945215.003.0208>.

45 Daniel J. Escott, "Verify or Perish: The New Standard of Professional Competence After *Ko v. Li*," *Internet and Technology Law in Canada* 26, no. 10 (February 2026).

46 Nathan Afilalo, Daniel Escott, and Archie Zariski, "Multi-Functional Access to Justice Centres," *Windsor Yearbook of Access to Justice* 40 (2024): 262-283; Daniel James Escott, *Meaningful Justice Design: A Practical Implementation of Accessible Justice* (LL.M. thesis, York University, July 2025).

adjudication are therefore useful high-consequence settings in which to test CSD's human-authority, evidentiary, review, and remedial dimensions. The broader proposition is narrower than a claim of substantive justice: technological delegation remains constitutionally defensible only if the responsible institution retains the capabilities required to exercise authority lawfully, independently, transparently, and subject to meaningful review and remedy.

LIMITATIONS

CSD has important limitations that should be stated before empirical testing. First, the term constitutional must be used carefully. In this research it refers to higher-order legal and institutional constraints on the legitimate exercise of public authority, including lawful authority, fundamental rights, procedural protections, accountability, review, and remedy. The Canadian Charter and Canadian administrative law provide an important implementation context, but the doctrine does not assume that every jurisdiction has the same constitutional structure or that identical institutional controls will be required in every legal system.

Second, the proposed dimensions may not be equally important across use cases. A low-impact administrative support tool may not require the same evidentiary or review architecture as a system influencing liberty, benefits, taxation, immigration, health, employment, or enforcement. The framework therefore requires contextual thresholds. Third, expert judgement may itself reflect disciplinary assumptions. Legal experts may emphasise authority and remedy, engineers may emphasise operational control and observability, security experts may emphasise resilience and access, and public administrators may emphasise feasibility and institutional capacity. The Delphi and reliability stages should treat those disagreements as data rather than noise.

Fourth, retained control does not establish substantive fairness. A government could exercise very strong control over a system that produces discriminatory, inaccessible, or otherwise unlawful outcomes. CSD is therefore not a complete measure of trustworthy AI, justice, or legitimacy. It identifies whether the institution retains the capabilities required to govern the system and answer for its decisions; those capabilities must still be exercised consistently with substantive law, rights, ethics, evidence, procedural fairness, and meaningful access to review.

Finally, the doctrine should not be used to justify technological isolation. Sovereignty and interoperability are not opposites. Canadian and European policy both acknowledge the practical necessity of global technological interdependence. The relevant question is whether dependencies are understood, governable, auditable, replaceable where necessary, and compatible with the institution's continuing ability to discharge public responsibilities. CSD therefore evaluates retained capability, not technological autarky.

CONCLUSION

The central argument of this research is that public accountability and technical control can no longer be assumed to occupy the same institutional or territorial space. AI-enabled public decisions may be assembled from data, infrastructure, models, services, and expertise

distributed across multiple organisations and jurisdictions. Yet the legal responsibility for the decision can remain concentrated in one public institution. This creates a governance problem that cannot be resolved through data residency, contractual allocation of responsibility, or nominal human oversight alone.

Constitutional Intelligence provides the institutional foundation for addressing this problem because it treats trustworthy AI as a decision architecture connecting authority, evidence, inference, human judgement, explanation, review, and remedy. Constitutional Sovereignty Doctrine extends that architecture through a principle of retained control: where legal and constitutional accountability remains with a public institution, the practical capabilities necessary to exercise that responsibility cannot be assumed to migrate harmlessly to external systems, suppliers, or actors. Accountability-Control Asymmetry identifies the condition that arises when responsibility remains but effective capability migrates. The Retained Accountability Principle provides the associated normative proposition that delegation should not extinguish the capacities required to inspect, contest, intervene in, reconstruct, verify, and remedy a public decision.

The doctrine also makes empirical claims that must be tested. Its eight dimensions should be examined for validity, reliability, discriminatory value, and practical usefulness. The assessment framework should be compared with simpler sovereignty proxies and revised if independent experts cannot apply the dimensions consistently, or if particular dimensions prove redundant, unobservable, or strongly context-dependent. CSD is deliberately open to empirical refinement; it does not require constitutional sovereignty to be reduced prematurely to a universal numerical index.

The argument can be stated simply: territory matters, but control matters more when digital systems are distributed; control, in turn, is insufficient when it is detached from the legal and constitutional obligations attached to public authority. Constitutional Sovereignty Doctrine joins those propositions by requiring responsibility for a consequential public decision to remain accompanied by sufficient practical control to make that responsibility meaningful.

REFERENCES

- Alon-Barkat, Saar, and Madalina Busuioc. "Human-AI Interactions in Public Sector Decision Making: 'Automation Bias' and 'Selective Adherence' to Algorithmic Advice." *Journal of Public Administration Research and Theory* 33, no. 1 (2023): 153-169. <https://doi.org/10.1093/jopart/muac007>.
- Afilalo, Nathan, Daniel Escott, and Archie Zariski. "Multi-Functional Access to Justice Centres." *Windsor Yearbook of Access to Justice* 40 (2024): 262-283.
- Baker v. Canada (Minister of Citizenship and Immigration), [1999] 2 S.C.R. 817.
- Canada (Minister of Citizenship and Immigration) v. Vavilov, 2019 SCC 65, [2019] 4 S.C.R. 653.
- Cecil, Julia, Eva Lerner, Matthias F. C. Hudecek, Jan Sauer, and Susanne Gaube. "Explainability Does Not Mitigate the Negative Impact of Incorrect AI Advice in a Personnel Selection Task." *Scientific Reports* 14 (2024): 9736. <https://doi.org/10.1038/s41598-024-60220-5>.
- Citron, Danielle Keats. "Technological Due Process." *Washington University Law Review* 85, no. 6 (2008): 1249-1313.

Council of Europe. Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law. Strasbourg: Council of Europe, 2024.

Dudoit, Alain, Tony Labillois, and Celio Oliveira. "Public Sector Data Interoperability and Responsible Adoption of AI in Canada: Strategic Overview and Call to Action." CIRANO 2026PR-02, 2026. <https://doi.org/10.54932/PUUN9151>.

Escott, Daniel J. "An Intertraditionalist Approach to 'Humans in the Loop.'" In *Oxford Intersections: AI in Society*, edited by Philipp Hacker. Oxford University Press, published 23 October 2025. <https://doi.org/10.1093/9780198945215.003.0208>.

Escott, Daniel J. "Verify or Perish: The New Standard of Professional Competence After *Ko v. Li*." *Internet and Technology Law in Canada* 26, no. 10 (February 2026).

Escott, Daniel James. *Meaningful Justice Design: A Practical Implementation of Accessible Justice*. LL.M. thesis, York University, July 2025.

European Commission. *Cloud Sovereignty Framework: Implementation Guidance*. Brussels: European Commission, 2026.

European Commission. "Commission Advances Cloud Sovereignty through Strategic Procurement." 17 April 2026.

European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence. *Official Journal of the European Union*, 2024.

Government of Canada. "Digital Sovereignty: A Framework to Improve Digital Readiness of the Government of Canada." 12 November 2025.

Government of Canada. "Government of Canada White Paper: Data Sovereignty and Public Cloud." 2018.

Kroll, Joshua A., Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. "Accountable Algorithms." *University of Pennsylvania Law Review* 165, no. 3 (2017): 633-705.

Mitchell, Andrew D., and Theodore Samlidis. "Cloud Services and Government Digital Sovereignty in Australia and Beyond." *International Journal of Law and Information Technology* 29, no. 4 (2021): 364-394. <https://doi.org/10.1093/ijlit/eaac003>.

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. Gaithersburg, MD: NIST, 2023. <https://doi.org/10.6028/NIST.AI.100-1>.

OECD. "OECD AI Principles." Paris: Organisation for Economic Co-operation and Development. Accessed 2 October 2026. <https://www.oecd.org/en/topics/ai-principles.html>.

Oliveira, Celio. "Constitutional Intelligence as a Governance Architecture for Trustworthy Public-Sector AI." *Open Access Journal of Data Science & Artificial Intelligence* 4, no. 1 (2026): 000167. <https://doi.org/10.23880/oajda-16000167>.

Oliveira, Celio. *Constitutional Intelligence: A Decision Architecture for Trustworthy AI Governance*. Altona, MB: FriesenPress, 2026.

Raji, Inioluwa Deborah, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. "Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33-44. New York: ACM, 2020. <https://doi.org/10.1145/3351095.3372873>.

Selbst, Andrew D., and Julia Powles. "Meaningful Information and the Right to Explanation." *International Data Privacy Law* 7, no. 4 (2017): 233-242. <https://doi.org/10.1093/idpl/ix022>.

Treasury Board of Canada Secretariat. "Algorithmic Impact Assessment." Open Government Portal. Accessed 2 October 2026. <https://open.canada.ca/aia-eia-js/?lang=en>.

Treasury Board of Canada Secretariat. "Directive on Automated Decision-Making." Modified 24 June 2025. <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592>.

Wang, Blair, Raffaele Fabio Ciriello, and Alexander Richter. "Digital Sovereignty." *Business & Information Systems Engineering* (2026). <https://doi.org/10.1007/s12599-026-01018-9>.

Wilkinson, Mark D., et al. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data* 3 (2016): 160018. <https://doi.org/10.1038/sdata.2016.18>.