

Bureaucratic Silences: What the Canadian AI Register Reveals, Omits, and Obscures

Dipto Das
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
dipto.das@utoronto.ca

Syed Ishtiaque Ahmed
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
ishtiaque@cs.toronto.edu

Christelle Tessono
Faculty of Information
University of Toronto
Toronto, Ontario, Canada
christelle.tessono@mail.utoronto.ca

Shion Guha
Faculty of Information
University of Toronto
Toronto, Ontario, Canada
shion.guha@utoronto.ca

Abstract

In November 2025, the Government of Canada operationalized its commitment to transparency by releasing its first Federal AI Register. In this paper, we argue that such registers are not neutral mirrors of government activity, but active instruments of **ontological design** that configure the boundaries of accountability. We analyzed the Register’s complete dataset of 409 systems using the Algorithmic Decision-Making Adapted for the Public Sector (ADMAPS) framework, combining quantitative mapping with deductive qualitative coding. Our findings reveal a sharp divergence between the rhetoric of “sovereign AI” and the reality of bureaucratic practice: while 86% of systems are deployed internally for efficiency, the Register systematically obscures the human discretion, training, and uncertainty management required to operate them. By privileging technical descriptions over sociotechnical context, the Register constructs an ontology of AI as “reliable tooling” rather than “contestable decision-making.” We conclude that without a shift in design, such transparency artifacts risk automating accountability into a performative compliance exercise, offering visibility without contestability.

CCS Concepts

• **Social and professional topics** → **Governmental regulations**;
• **Applied computing** → **E-government**; • **Human-centered computing** → Empirical studies in collaborative and social computing.

Keywords

AI Register, Government of Canada, Bureaucracy, Accountability, Transparency

ACM Reference Format:

Dipto Das, Christelle Tessono, Syed Ishtiaque Ahmed, and Shion Guha. 2026. Bureaucratic Silences: What the Canadian AI Register Reveals, Omits,

and Obscures. In *The 2026 ACM Conference on Fairness, Accountability, and Transparency (FAccT '26)*, June 25–28, 2026, Montreal, QC, Canada. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3805689.3812336>

1 Introduction

In November 2025, the Government of Canada (GC) publicly released its first government-wide artificial intelligence (AI) Register, disclosing over 400 AI systems used across more than 40 federal institutions [125]. The GC highlighted making “public service more efficient” as the motivation for incorporating AI into federal government operations, for which this register would reduce duplication, and help departments identify opportunities to work more efficiently [125]. Various geopolitical entities are building similar registers [49, 55, 79, 134], such as the European Union (EU)’s Digital Services Act (DSA), positioning such disclosure infrastructures as a central mechanism for governing algorithmic systems used in the public sector. From the algorithmic fairness, accountability, and transparency (FAccT) research perspective, we ask: *how meaningfully do AI registers enable accountability, or do they merely institutionalize its appearance?*

FAccT research has largely framed disclosure, through documentation standards, audits, and reporting requirements, as a necessary condition for oversight [95]. While these approaches have advanced important norms (e.g., datasheet, model cards) [45, 76], critical scholarship [116] has shown that institutional factors (e.g., incentives, structures) substantially shape disclosure efforts, such as AI registers. Moreover, emerging discourse on AI governance often treats AI systems as conceptually interchangeable, increasingly conflating generative and large language model (LLM)–based systems [127]. This emphasis on highly visible and popular tools obscures the continued operation of long-standing predictive, scoring, and decision-support systems that continue to shape public-sector decision-making [140]. By rendering these systems visible and classifying them in particular ways, AI registers shape how accountability is practiced, and more significantly, define what counts as AI in governance practice.

Rather than asking what information the Canadian AI Register discloses, we examine how it delineates responsibility, represents uncertainty, and locates discretion. To do so, we used the algorithmic decision-making adapted for the public sector (ADMAPS)



This work is licensed under a Creative Commons Attribution 4.0 International License. *FAccT '26, Montreal, QC, Canada*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2596-8/2026/06
<https://doi.org/10.1145/3805689.3812336>

framework [111], which foregrounds various types of interactions within public-sector governance. Thus, this paper makes three contributions. First, we provide an empirical analysis of the complete Canadian federal Public AI Register to systematically examine 409 documented systems. Second, we introduce the concept of *bureaucratic silences* to describe how AI registers structure what is disclosed about public sector AI systems (e.g., tool- and developer-level disclosures, technical capabilities, and efficiency), and what remains illegible or off the record (e.g., human, bureaucratic, and contextual factors), particularly around discretion, infrastructure, and uncertainty. As similar registers emerge across Europe and beyond [49, 55, 79], the silences we identify and problematize should be examined as systemic properties of register-driven transparency regimes, rather than uniquely Canada-specific peculiarities—an insight that broadens the relevance of our analysis for global AI governance. Third, we argue that AI registers should be understood as instruments of *ontological design*, which shape how accountability is defined and enacted, with concrete implications for how future registers should be designed if they are to support meaningful democratic oversight and public trust in public-sector AI. We situate our findings by discussing how the AI register contributes to and articulates Canada’s preparedness for its national AI governance and sovereignty strategies.

2 Literature Review

Our literature review brings together three bodies of work that frame how algorithmic accountability is enacted in practice: sociotechnical perspectives on accountability and transparency, research on documentation and disclosure as governance infrastructures, and critical studies of public-sector algorithmic systems. We draw on these strands to motivate our choice of ADMAPS as the analytical lens for examining the Government of Canada (GC) AI Register.

2.1 Algorithmic Accountability and Transparency in Sociotechnical Systems

Accountability and transparency have emerged as central concepts in scholarship on the social impacts of algorithmic and AI systems. Whereas social sciences understand accountability as the everyday person’s ability to hold an institution responsible [64], from a technical standpoint, accountability refers to the capacity of algorithmic systems to trace, justify, and assign responsibility for their decisions and outcomes [14, 36, 136], whereas transparency is commonly understood as the availability of information that enables explanation, documentation, or disclosure of system properties [12, 63, 131]. These concepts are often operationalized through technical and procedural interventions, such as explainable models, documentation frameworks, audit mechanisms, and reporting standards, which aim to render algorithmic systems more legible to developers, regulators, and affected users [43, 76, 123].

Beyond such technical understandings of algorithmic accountability and transparency, scholars increasingly conceptualize them as shaped by sociocultural contexts, institutional logics, organizational structures, legal mandates, political priorities, and power asymmetries [2, 30, 116, 126]. Within this line of critique, accountability is reframed as a question of governance: who is accountable

to whom, through what mechanisms, and under what institutional conditions [14, 136]. For example, studies have shown that algorithmic supply chains fragment responsibility across vendors, developers, and deploying institutions, complicating traditional notions of oversight and liability [26]. In such a scenario, where responsibility for algorithmic decisions is distributed across organizational roles, infrastructures, and decision points rather than residing solely in models or code, traceability emerges as a principle for operationalizing accountability [60]. In this governance-oriented view, transparency is not reducible to the disclosure of technical details but is better understood as a “communicative constellation” and a form of “disclosure by design,” in which disclosures are situated, audience-specific, and shaped by institutional incentives, actively structuring how information is interpreted, which forms of accountability become possible, and who is empowered to act [40, 81].

Critical scholars have further highlighted the organizational stakes of accountability and transparency practices themselves, such as how disclosure datasets embody politics in defining what counts as accountable knowledge [16, 98]. They showed how accountability is often pursued through collective action (e.g., advocacy) as a negotiation between efforts to strengthen accountability through disclosure and attempts to reduce regulatory burden. Algorithmic accountability depends on alignment among system design, professional judgment, and organizational context [132], and is an ongoing practice rather than a one-time technical intervention [53]. Redden and colleagues have examined its extent in the public sector through case studies in fraud detection, child welfare, social services, and policing [31, 101, 103] across Europe, North America, and Australia [102]. In the Canadian context, Redden [100] similarly investigated government discourses and practices surrounding big data adoption in the public sector using counter-mapping methods and freedom-of-information requests. These studies demonstrated significant gaps in existing transparency regimes since algorithmic governance is often difficult to trace due to fragmented documentation practices, institutional opacity, and limited proactive disclosure. Overall, accountability is a relational and institutional practice that requires examining the documentation norms, disclosure regimes, and governance artifacts. Hence, researchers have called for systematic documentation and public disclosure mechanisms that enable oversight of algorithmic systems in governance.

2.2 Accountability in Governance through Documentation and Disclosure

Sociotechnically grounded AI governance research highlights how documentation and disclosure, operating at the intersection of technical practice and organizational governance, translate abstract accountability commitments into standardized, inspectable artifacts that circulate across organizations, regulators, and publics [45, 92, 138]. Examples of these artifacts are documentation frameworks such as datasheets and model cards, which promote responsible AI development by standardizing disclosures about provenance, intended use, limitations, and risks. These frameworks also function as community practices, embedding disciplinary norms, institutional incentives, and implicit assumptions about responsibility and audience into their design and use [34, 72]. However, documentation does not simply describe underlying systems—instead, it

actively shapes which forms of information are legible and what kinds of accountability claims become possible. Therefore, data systems' accompanying documentation must be "read" critically: what decisions they encode about classification and how those would shape downstream interpretation and governance [61, 97].

Documentation practices not only support transparency and accountability but also delimit their scope by foregrounding certain risks, actors, and values while understating others. Claims about scale, representativeness, or neutrality often become central to these documentations and obscure the contingent social and organizational conditions under which data are produced and maintained [58]. For instance, scholars have traced how the discourse around data systems stabilizes narratives [33] by shaping how influential benchmarks are conceptualized, reducing concerns about fairness and accountability to concerns about sufficient data [3, 52], and by making certain types of data labor invisible [115, 117].

Beyond individual datasets or models, normative demands for accountability and transparency are often enacted through formalized governance infrastructures and structured reporting obligations. AI registers have emerged as a governance instrument in several jurisdictions seeking to increase transparency around public-sector AI systems. Municipal governments such as Amsterdam [79] and Helsinki [49] were among the earliest adopters of public algorithm registers, publishing structured documentation about the AI systems used in public services, which have since inspired similar calls for disclosure mechanisms in other jurisdictions, such as the EU's DSA [59] Bangladesh's AI policy [82], and Canada's municipal and organizational AI registers [118, 120]. Empirical studies focusing on these transparency databases, AI registers, ethical charters, and compliance repositories across various levels, scales, and contexts have examined how documenting patterns in self-reported actions, inconsistencies in categorization, and limits to contextual explanation [38, 56, 130] standardize disclosure pipelines and moderation practices that prioritize structured, machine-readable reporting over interpretive openness [55]. Scoping these insights in the context of the public sector, researchers have analyzed how disclosure templates, reporting thresholds, and organizational incentives shape what institutions can or choose to document about deployed systems [51, 80], similar to Redden et al.'s insights [100, 102] we discussed earlier. Because of these design choices, accountability initiatives often fall short of their regulatory objectives [48]. These challenges often exacerbate in Global South contexts due to different AI ethics that underpin these policies [13, 69, 93].

Related works on AI ethics charters and manifestos map how normative voluntary disclosures and commitments across institutions produce an appearance of ethical accountability at a meso-governance level, without necessarily mandating enforceable accountability [27, 46, 50]. Taken together, this body of work traces a shift from transparency as an aspirational principle toward transparency as an infrastructural practice. Similarly, studies of large-scale data infrastructures, like Common Crawl, demonstrate how disclosure practices can obscure responsibility by diffusing accountability across actors, licenses, and downstream users [10]. In the context of fairness and accountability research, such "lazy" data documentation practices can entrench methodological shortcuts and epistemic blind spots, reinforcing the appearance of accountability without substantively addressing underlying harms [122].

2.3 Critical Analysis of Accountability of Algorithmic Systems in the Public Sector

Algorithmic accountability studies employ a wide set of empirical approaches. A prominent method among these is auditing, which examines decision-making patterns while varying inputs to surface bias, error, or harm [99, 107]. Both researchers and everyday users have used algorithmic audits to scrutinize AI systems [32, 73, 119], like recommendation and search algorithms [11, 106], computer vision-based processes (e.g., generative art [124], image captioning [139], facial recognition [24]), and language technologies (e.g., sentiment analysis [28, 57], hate-speech detector [108], machine translation [109], text generation [42]). Scholars have also emphasized that deciphering the algorithmic outputs alone is often insufficient for rigorous accountability analysis [44], particularly when system behavior is shaped by organizational workflows, data pipelines, and governance constraints that remain undocumented or inaccessible [25, 70, 112]. In such cases, investigative and journalistic approaches, combining technical inspection with document analysis, interviews, and policy tracing, are also effective in identifying opaque practices and institutional accountability gaps [35, 37].

While impact assessments structure accountability by anticipating and documenting potential harms, trade-offs, and mitigation strategies, their effectiveness depends on institutional incentives and implementation practices [135]. In public services, accountability challenges are amplified by high stakes, legal mandates, and complex procurement arrangements [22]. To navigate through these, public-sector algorithms often redistribute discretion rather than eliminate it, shifting responsibility across frontline workers, technical systems, and administrative hierarchies [78, 113, 132]. In the Canadian context, the Directive on Automated Decision-Making [85], introduced by the Treasury Board Secretariat in 2019, shapes such algorithmic governance practices. The directive requires federal agencies deploying automated decision systems to complete an Algorithmic Impact Assessment (AIA) [91] that evaluates potential impacts of algorithmic decision-making on rights, fairness, and administrative accountability. While the AIA functions primarily as an ex-ante assessment mechanism, the GC AI Register serves as a disclosure infrastructure, which documents and describes systems in development or already deployed across federal agencies.

To critically examine the GC AI Register, we draw on the algorithmic decision-making adapted for the public sector (ADMAPS) framework [111]. Because it focuses on how algorithmic systems are embedded within bureaucratic processes, institutional practices, and data infrastructures rather than treating them solely as technical artifacts, the ADMAPS framework offers a useful analytical lens for our study. Unlike actor/stage mappings that catalog governance instruments [61] or tools for AI impact assessment in terms of prospective risks [54], ADMAPS offers an analytical lens for reading the GC AI Register as an accountability artifact embedded in the public sector. This framework conceptualizes algorithmic governance as a sociotechnical arrangement shaped by the interaction of human discretion, bureaucratic processes, and algorithmic decision-making. Human discretion encompasses professional expertise, value judgments, and heuristic decision-making exercised

by frontline practitioners. Bureaucratic processes capture organizational conditions such as resources and constraints, administrative routines and training, and governing laws and policies. Algorithmic decision-making attends to the role of data, the type of decision support provided (e.g., predictive or prescriptive), and the degree of uncertainty embedded in algorithmic systems. Overall, this offers a structured lens for analyzing how the GC AI Register defines, represents, organizes, and assesses the AI systems across various departments and agencies of the Canadian government.

3 Methods

In this section, we explain our mixed-methods approach to analyzing how AI systems are framed in the GC AI Register, combining mixed-method exploration with ADMAPS-guided critical discourse analysis.

3.1 Data Collection, Preprocessing, and Exploratory Analysis

Our study draws on the GC AI Register [88], which provides structured, inventory-like information on 409 AI systems, ranging from early initiatives to fully operational tools supporting service delivery and internal operations across 42 federal departments and agencies. Each entry includes fields, such as: ID, AI system's name, description, primary users, status, date, capabilities, data sources, results, as well as which government organization is associated with it, who developed it, and vendor information, and whether and which personal information banks are involved with it, in both English and French. Our analysis focuses on how public-sector AI systems are represented and documented in the register, rather than on auditing the registered AI systems or offering an ethnographic account of their day-to-day use. In other words, our object of study is the register as a sociotechnical governance artifact that projects particular views of data, decision-making, and accountability, in line with the ADMAPS dimensions. In this paper, we treat the register not only as an administrative dataset but also as an institutional governance artifact that reflects how AI use is documented, framed, and communicated in the context of public-sector policy commitments and accountability norms.

In the tabular dataset, each row corresponds to a unique AI system. We retained only the English columns and added a column for the acronyms of the corresponding government organizations, as listed in [87]. We formatted the categorical fields for consistency (e.g., capitalization and spacing) and imputed empty categorical fields as 'Unspecified'. We exploratorily analyzed the quantitative distributions of nominal variables: primary users, associated government organization, development status, vendor information, and the use of personal information banks. To examine patterns in reported AI capabilities, we applied lightweight text analysis to the capability descriptions after removing stop words and common domain-generic terms (e.g., 'model,' 'data,' 'AI'), extracting tokens, bigrams, and trigrams, and generating a word cloud to surface recurrent capability themes across systems in the register.

3.2 Data Analysis

We used deductive qualitative analysis guided by the ADMAPS framework [111]. Using it as the analytical lens, the first author

conducted the initial open coding, i.e., identified recurring representational patterns: entities, tasks, objectives, etc., that frequently appeared in the textual descriptions of AI systems and their results within predefined theoretical dimensions. Using each AI system entry in the Register as the unit of analysis, the codebook was iteratively developed: the codes generated by the first author were refined through discussions with the other three authors, where disagreements or ambiguities were documented in analytic memos and collaboratively resolved, and codes were merged, split, or adjusted to maintain consistency. Following social computing research guidelines for such approaches [71], we did not calculate the inter-coder reliability score for the codes. To examine how AI systems are framed, we mapped the codes in relation to the ADMAPS dimensions: algorithmic decision-making, bureaucratic procedure, and human discretion [111], explained earlier in Section 2. Drawing on critical discourse analysis that considers what is said, emphasized, and left unsaid in texts and documents [41, 74], we examine the register as an institutional governance artifact whose standardized descriptions and categories structure how AI systems and accountability relations are represented. Accordingly, we treated both recurring emphases and systematic silences as analytically meaningful in our iterative coding. For example, while coding for degrees of uncertainty, we noted explicit references to pilots or experimentation, as well as the absence of such references in descriptions of mature or operational systems. We juxtaposed exploratory quantitative insights with qualitative interpretation (e.g., considering frequency as an indicator of institutional normalization) to understand how public-sector AI systems are framed within federal governance.

3.3 Limitations

This study has several limitations. First, our analysis is bounded by what is documented in the register and cannot access undocumented systems, informal workarounds, or the situated practices of frontline workers. Hence, our study illuminates the politics of disclosure and representation rather than the operational performance or impacts of each system. Second, our coding approach is interpretive and theory-driven, anchored in ADMAPS. While other frameworks could foreground different aspects of the same entries, our objective is to critically examine the Canadian federal AI systems with reference to their bureaucratic situatedness and institutional logics.

4 Results

We organize our results through the ADMAPS framework, examining (1) how algorithmic decision-making is represented and operationalized; (2) how AI systems are described in relation to bureaucratic processes; and (3) how human discretion, expertise, and heuristics are reconfigured. Across each category, we examined its corresponding elements.

4.1 Algorithmic Decision-Making

In this section, we analyze how algorithmic decision-making is represented in the GC AI register by identifying the type of data, the forms of decision-making support, and the degrees of uncertainty embedded in the algorithmic systems.

decision-makers, such systems can narrow interpretive latitude by nudging workers toward preferred pathways, particularly when aligned with performance metrics.

We also identified systems that align with ADMAPS's definition of **predictive systems** that estimate risks, forecast future outcomes, or estimate likelihoods. For instance, the CBSA's Courier Low Value Shipments uses "probabilistic models for the risk scoring of packages", while ESDC's Employment Insurance ML Workload works on "identifying cases where a recalculation will not result in any change.". Similar dynamics appear in metric-driven systems at Environment and Climate Change Canada and Fisheries and Oceans Canada, such as automated oil spill detection and oceanographic anomaly detection. By formalizing uncertainty into decision pipelines, predictive systems reshape where discretion accumulates and may encourage officers to prioritize model-flagged cases over situational judgment, a form of algorithmic over-reliance identified in ADMAPS.

Finally, several systems **automate operational tasks** to optimize workflows and delegate procedural work. For example, ECCC's OCR/LLM CITES Permit Project converts scanned documents into structured data, while agencies such as Natural Sciences and Engineering Research Canada use AI to automate data entry, detect and duplicate documents, and support grant analysis. By mechanizing tasks once handled by experienced staff, operational automation reshapes the boundary between routine and discretionary work, thus boosting efficiency but potentially eroding opportunities for tacit knowledge development that ADMAPS identifies as central to professional expertise. When automation reorganizes workload distribution, it also restructures evaluative horizons: workers may be judged on throughput rather than judgment, redefining professional identity and expectations.

4.1.3 Degrees of Uncertainty. Examining the Canadian AI Register entries through the lens of ADMAPS reveals a spectrum of uncertainty expressions—ranging from explicit acknowledgment, to strategic minimization, to complete omission. These patterned articulations of uncertainty shed light on how agencies anticipate human discretion, envision bureaucratic processes, and justify the deployment of algorithmic tools.

A large subset of systems **explicitly acknowledge uncertainty** through linguistic framing like "pilot," "experiment," "testing," "exploring," "intended to establish infrastructure," or "in development". This positioning treats algorithmic models as provisional and exploratory, reflecting ADMAPS's insight that uncertainty should be made visible so that their primary users can calibrate their interventions. For example, systems such as Westlaw Edge with CoCounsel, Lexis+ AI, or Nuix Neo are openly framed as "pilot only," with evaluations underway comparing AI-generated outputs to conventional human workflows. Other systems, such as Risk Scoring for CLVS, APIDS, or AI-assisted weather prediction, reflect modeling uncertainty, requiring generalization from noisy or incomplete data. Moreover, the Administrative Tribunals Support Service of Canada describes an AI system as "the required infrastructure, tool, and skill foundation" for future AI and definitive systems to emerge. Their descriptions often include qualifiers such as "may allow," "exploring," or "intended to streamline," foregrounding that, in ADMAPS

terms, their "degrees of uncertainty" in outcomes remain large and algorithmic performance cannot yet be relied upon as authoritative.

In contrast, other entries **present AI systems as stable** by framing those as mature, standardized, or evidence-driven, emphasizing improvements in "accuracy," "reliability," "efficiency," or "consistency." Systems such as Document Detective, Electronic SIN Automation, Record of Employment Comments Assessment, or Export Compliance Dashboard use language that implies performance stability, even when the underlying tasks (e.g., identity verification, anomaly detection, document comparison) are prone to classification error and contextual ambiguity. For instance, Electronic Social Insurance Number (SIN) Automation highlights its ability to verify identity documents using computer vision, suggesting near-real-time reliability. Using a blanket phrase like "when certain criteria are met," the entry tries to downplay the inherent uncertainty in document authentication, due to variations in image quality, emerging fraud vectors, and contextual cues. By foregrounding efficiency over limitations, these systems shift uncertainty away from the algorithm and onto human officers who must resolve edge cases after confident classifications are produced. ADMAPS warns that such a performative impression of stability risks narrowing the perceived space for human discretion. Thus, uncertainty is not eliminated, but is obscured.

A third category consists of systems whose descriptions contain **no explicit reference to uncertainty**. They reflect on no risks, no limitations, no error surfaces. This silence itself is meaningful. Many systems describe tasks involving classification, summarization, or clustering—activities that inherently entail sociotechnical uncertainty. Still, their descriptions focus solely on productivity, speed, or convenience, offering no guidance on how human discretion should be configured in light of algorithmic fallibility. This silent omission effectively treats uncertainty as external to the system and relocates responsibility to human actors without naming it as such.

4.2 Bureaucratic Processes

We examined how the bureaucratic processes surrounding AI systems are shaped by GC's internal and external resources and constraints, administrative scaffolds and training expectations, as well as federal and provincial laws and policies.

4.2.1 Resources and Constraints. We examined how the intended results of the AI systems are framed in the register. Many systems explicitly emphasize goals such as "reducing workload," "managing backlogs," "improving efficiency," or enabling cases to be "attended to quickly". These recurring justifications position algorithmic systems as compensatory infrastructural responses to staffing shortages, procedural burdens, and chronic workload pressure. ADMAPS conceptualizes such constraints and related resources as material conditions (e.g., staffing, expertise, funding, and organizational bandwidth) that shape how algorithmic systems are adopted and integrated into bureaucratic practice.

There are over 200 departments and agencies in the Government of Canada [87], yet only 42 organizations appear on the AI register. Of which, an even smaller number of organizations, such as the Canada School of Public Service (CSPS), Innovation, Science and Economic Development Canada (ISED), National Research Council

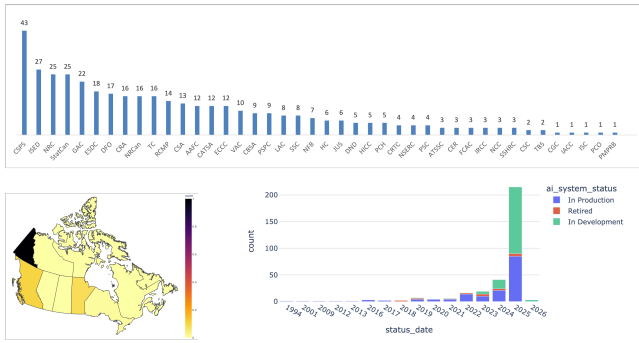


Figure 2: Top: Distribution of AI systems across GC organizations; Bottom-left: Provinces and Territories explicitly mentioned in the AI register using population-normalized values; Bottom-right: Distribution of AI systems across years and status.

Canada (NRC), Statistics Canada, GAC, and ESDC, account for the majority of systems (see Figure 2: top). Since 2022, the register shows a temporal trend of increased AI development and deployment (see Figure 2: bottom-right): 44% of systems are “In Development,” while 39% are already “In Production,” implying simultaneous expansion and maintenance rather than isolated projects. We also observed institutional borrowing and infrastructural reuse, where systems developed within one organization inform deployments in other departments and agencies. For example, Justice Canada (JUS)’s adoption of Westlaw Edge mirrored evaluations of Lexis+ AI and Nuix Neo, while GAC’s partnership with NRC enabled early deployment of scalable NLP-based solutions. ADMAPS notes that continuous system building under such conditions can limit opportunities for reflection, retraining, or recalibration of bureaucratic processes.

Patterns in system development strategies signal agencies’ resource availability and their navigation of capacity constraints (see Figure 3). Systems are almost evenly split between in-house development (43.3%) and external vendors (38.1%), while open-source development remains limited (6.1%). A small number of vendors, most notably Microsoft (n=46), account for a disproportionate share of externally developed systems. This concentrated reliance on a small number of foreign technology providers in public-sector procurement dynamics can shape both the technological choices and governance practices surrounding public-sector AI systems. Moreover, we identified three development strategies. First, exclusive in-house development (e.g., National Defence, Housing, Infrastructure and Communities Canada), indicating sustained internal technical expertise. Second, full reliance on external vendors or the adoption of open-source tools. And third, a mixed approach with some adopting a fully third-party approach, sometimes depending entirely on a single external vendor (e.g., Darktrace, Microsoft) or primarily on open-source systems. A third group follows a mixed strategy, combining in-house development with external vendors. For example, CSPS draws on a particularly large and heterogeneous vendor ecosystem involving 33 different entities. Under the ADMAPS framework, these differences are not merely technical

choices but reflect how agencies allocate scarce resources, navigate procurement constraints, and balance short-term operational demands against longer-term capacity development.




Figure 3: Proportion of developers by GC organizations.

4.2.2 Administration and Training. Administration and training foreground how organizational protocols, workflows, and capacity-building practices condition the operation of algorithmic systems in the public sector. Rather than treating AI systems as standalone technical artifacts, this dimension draws attention to the bureaucratic infrastructures (e.g., procedures, routinized practices, and institutional learning) that enable or constrain their use in practice.

Across the register, many systems are framed as **internal workflow instruments** rather than public-facing applications. For instance, these include service management chatbots (e.g., Milo 2.0), AI-assisted legal research platforms (e.g., Westlaw Edge), and internal generative AI tools for drafting, coding, summarization, or analysis (e.g., Genni). In these cases, AI is positioned as a means of optimizing bureaucratic throughput by *reducing wait times, accelerating internal workflows, automating repetitive tasks, improving quality, or reallocating staff attention toward “more complex” tasks*. Despite frequent references to the need for **adjusted procedures and workflows**, entries that mention systems integrating with ticketing tools, legacy databases, document repositories, or cross-departmental infrastructure often leave the new intake, handoff, and review protocols they implicitly require unexplained.

Similarly, **training requirements** are rarely articulated. Many AI systems entail layered, multi-step interactional workflows: threshold configuration, probabilistic interpretation, large-scale document retrieval, validation of AI-generated outputs, and coordination of human-in-the-loop review [1]. AI Systems (e.g., for medical adjudication or security screening) require both expertise and familiarity with AI interfaces and limitations, yet the register rarely specifies how workers are trained to interpret outputs, recognize failures, or integrate AI with professional judgment. Training is treated as implicit rather than a formal element of algorithmic governance. Even in pilot deployments, the focus remains on anticipated outcomes rather than on institutional preparation.

The GC AI register consistently presents AI as functional solutions, privileging a tool-centric representation of public-sector AI view that treats administrative processes and training infrastructures as peripheral or self-evident. This framing obscures the role of bureaucratic capacity and presents AI adoption as a technical intervention rather than an ongoing organizational transformation.

4.2.3 Laws and Policies. Laws and policies define formal constraints under which public-sector algorithmic systems are conceived, justified, and operationalized, encompassing not only legislation and

regulatory mandates but also jurisdictional boundaries that allocate authority. Examining the register through the ADMAPS framework reveals that references to provinces and territories are sparse and uneven: only a subset of jurisdictions are explicitly named (see Figure 2). The provincial or territorial references tend to appear in systems that must explicitly navigate multi-jurisdictional legal frameworks, such as regulatory compliance, enforcement, or coordination across federal–provincial–territorial boundaries. For example, tools (e.g., BizPaL) that aim to help their users navigate permits and licences across federal, provincial, and territorial regimes make jurisdiction explicit precisely because regulatory authority is fragmented. Here, AI is positioned as a mediator between overlapping legal systems, translating heterogeneous legislative requirements into standardized guidance. Conversely, the majority of AI systems in the register, particularly internal productivity tools, enterprise chatbots, and generative AI platforms, are described in jurisdiction-agnostic terms. Many of these systems are certified to process Protected B information (a classification under the GC’s security framework for handling medium-sensitivity government information [75]).

These systems operate within federal administrative authority, where legal constraints (e.g., privacy classification, information security, official languages requirements) are assumed to be uniform across the federal bureaucracy. This assumption becomes more complex for systems (e.g., Fuzzy Search) that rely on personal data banks while being developed or maintained by external vendors, many of whom are foreign entities. While Canada’s Personal Information Protection and Electronic Documents Act [84] governs private-sector processing of personal information in commercial contexts, the register provides little information on how privacy and data protection are federally governed in public-sector AI systems that rely on vendor-provided infrastructure.

In Canada, federal responsibilities include security and national defense, whereas provinces govern domains such as health, education, and welfare, with areas like immigration and agriculture falling under concurrent jurisdiction [121]. Though the register lists AI systems operating across several of these areas, it offers little visibility into how systems coordinate across jurisdictional domains. Instead, federal authority appears naturalized as the default governance layer, with provincial variation surfacing only when it introduces explicit regulatory complexity. The selective mentions of provinces and territories also underscore how law shapes what must be documented in the register. While most AI systems are implicitly framed as nationally applicable, the provincial specificity of some thus signals legal complexity rather than localized deployment. This pattern reinforces a tool-centric view of legality, in which legal complexity is acknowledged only when it cannot be abstracted away. The register does not, for instance, describe how federal AI systems interact with provincial privacy regimes, First Nations governance structures, or regional interpretations of shared mandates, despite these being consequential legal realities in practice.

4.3 Human Discretion

About 86.3% of AI systems in the register are explicitly designed for GC employees, with only a small share oriented toward public or

mixed-audience applications, highlighting the centrality of human value, heuristics, and expertise.

4.3.1 Professional Expertise. Following ADMAPS, we examine professional expertise as tacit, experiential, and domain-specific knowledge used by public-sector practitioners to interpret policy, evidence, and cases. Across the AI register, decision-making authority remains with trained practitioners, who are expected to interpret and act on system outputs. For example, legal research and drafting tools (e.g., Westlaw Edge AI, Lexis+ AI, and Otto) are framed as assistants for JUS lawyers, while Security Screening Automation supports officers by flagging areas of concern, and final decisions explicitly remain human-led. Our analysis examines how such systems augment professional judgment and whether expertise is standardized or formalized through the system itself. Many systems are framed as mechanisms to **augment professional expertise** by reducing cognitive and administrative burdens. Generative AI tools for summarization, search, translation, and coding, such as ProteBee, Document Detective, and Science-GPT, allow practitioners to reallocate time away from clerical labor toward interpretive or discretionary work. Notably, several systems explicitly emphasize capacity-building rather than automation: ProteBee aims at “*reducing the analysis processing time by at least 50%*”, while Science-GPT is co-developed with subject matter experts to limit hallucination and ground outputs in domain knowledge, positioning AI as a scaffold for professional judgment rather than a substitute.

At the same time, many systems **encode professional expertise into structured workflows**, predefined criteria, or assessment frameworks. Risk-scoring tools, triage systems, and classification models, such as Fuzzy Search and Travelers Compliance Indicators, translate historically tacit judgments into formalized rules, features, and thresholds. While such systems promise consistency, they also shift how expertise is expressed: from situated, case-by-case reasoning to standardized representations of what matters. This encoding is particularly visible in evaluation and policy contexts. Tools like GAC’s Enhanced Reporting for International Assistance explicitly embed the domain knowledge of experienced evaluators into machine-readable policy markers, producing outputs comparable to expert assessments while enabling oversight through explanations, making those structured, repeatable, and scalable across cases. However, one notable pattern across the registry entries is the limited presence of completed Algorithmic Impact Assessments, suggesting a gap between Canada’s formal assessment requirements under the Directive on Automated Decision-Making [85] and the systems publicly documented in the AI register.

4.3.2 Value Judgments. We found that the key values of efficiency/speed, accuracy/quality, equity/fairness/accessibility, and risk/safety in high-stakes contexts are all deeply embedded in both the design and application of AI systems in the register. For example, the Privy Council Office’s DataMinr First Alert is designed to enhance **efficiency** by summarizing large volumes of data in real time, thereby speeding up decision-making. Efficiency is often prioritized in time-sensitive contexts. However, prioritizing speed can undermine **accuracy**, a tension that systems like DeepL Pro and GitHub Copilot address by emphasizing error reduction in translation and code development—though inconsistent data practices raise questions

about the reliability of these gains. For example, First Alert’s ReGenAI feature relies on “*publicly available data sources accessed through the internet*”. Given that online data is often tainted by communal prejudices and stereotypes [94], concerns about bias and fairness arise. In particular, public-sector AI systems are expected to uphold principles of **equity, fairness, and accessibility**. For instance, Ombot was developed to enhance access to Ombuds office services while ensuring privacy and reducing barriers for employees seeking assistance. Similarly, ESDC’s EquiVision aims to improve adherence to Diversity, Equity, and Inclusion (DEI) commitments by automating the analysis of DEI data and making it easier for decision-makers to address employment equity issues. Finally, values regarding **risk and safety** are prominent in AI applications for regulatory and compliance purposes. Systems such as ISED’s Entity and User Behavior Analytics use machine learning to monitor activity patterns and detect anomalies, helping identify potential cybersecurity threats or insider risks. Here, the judgment centers on balancing the safety and security with the potential risk and ethical concerns of over-surveillance and privacy violations.

4.3.3 Heuristic Decision-Making. Heuristic decision-making refers to the rules of thumb, shortcuts, and experiential judgments that practitioners rely on when operating under uncertainty, time pressure, and resource constraints. Considering that, in public-sector contexts, such heuristics are not signs of irrationality but pragmatic adaptations to complex environments where exhaustive analysis is infeasible [68], we focus on how algorithmic tools become part of mechanisms such as triage, prioritization, and flagging. First, a large subset of systems explicitly frames their role as **helping staff prioritize** work. Such triaging mirrors common human heuristics, such as starting with the highest-risk cases, while formalizing them into algorithmic criteria. Analysts still decide whether to escalate a case, initiate an investigation, or dismiss a flag, but their discretion is influenced by how the space of possible actions is shaped in advance with algorithmically generated signals. Second, some systems frame their role as **supporting complex decisions** rather than filtering cases. Here, algorithmic heuristics include looking for precedents, suggesting relevant sections, and comparing with known patterns. These systems can be understood as externalizing experienced practitioner heuristics, i.e., making tacit strategies explicit, repeatable, and faster to apply. While this can reduce cognitive load and improve consistency, it also subtly reshapes how practitioners reason through cases. Third, many AI systems on the register describe their role as **flagging** instances based on anomalies or patterns. However, as discussed before, these systems do not decide outcomes; instead, they decide what deserves attention. For example, CBSA’s Fuzzy Search/SSAName3, which we discussed earlier in 4.1.1, generates ranked matches and security flags. This is a high-risk system with a high impact on people, yet the register is nearly silent about how officers review or interpret these signals in practice. This omission leaves unclear whether such flags function as advisory cues, mandatory screening triggers, or informal heuristics guiding discretionary border enforcement decisions. Such redistribution of attention is a powerful form of heuristic intervention, as it influences which cases receive scrutiny and which pass unnoticed. Overall, these interventions relocate heuristics

from the point of decision or discretion to upstream moments of attention-setting.

5 Discussion

This paper speaks to current and proposed AI/algorithm registers (e.g., in European cities and transnational contexts [49, 55, 79]). Our results suggest that the Canadian AI Register enacts a technical model of algorithmic accountability, emphasizing systems’ capabilities, efficiency gains, and functional outcomes while leaving administrative processes, training arrangements, and institutional capacity largely unclear. This framing presents the GC’s adoption of AI as a technical choice or solution rather than an organizational transformation, which obscures the bureaucratic processes necessary to use these technologies. In terms of uncertainty, some descriptions convey experimentation, others convey stability and provisionality, while many omit discussion of it entirely. This forces frontline practitioners to infer failure modes and manage ambiguity, thus constraining discretion and increasing the system’s practical burden as these tools shape upstream triage and attention. In the following section, we discuss how this study’s findings align with Canada’s approach to AI governance, with research within the broader FAcCT community, and with public trust in governance.

5.1 Implication for Canadian AI Governance

Canada organizes AI governance through a bifurcated oversight model that differentiates between public-sector and private-sector AI operations. This has resulted in a regime characterized by substantial federal investments in private-sector AI development, alongside weak and decentralized regulatory oversight frameworks [6, 20, 128]. There are no centralized laws, regulations, or frameworks; instead, a patchwork of codes, guidelines, and directives is unevenly applied to public and private sector institutions. Flagship initiatives such as the Pan-Canadian AI Strategy and the Canadian Sovereign AI Compute Strategy have prioritized research capacity, commercialization, standards development, talent attraction, significant investments in compute infrastructure, and institutional support [83, 86], positioning Canada as a global AI hub. However, their outcomes raised concerns of disproportionately benefiting industry actors, offering limited mechanisms for public accountability, and externalizing public benefits to foreign actors [17, 105].

Although public-sector AI systems are formally governed by Treasury Board Secretariat policy instruments, such as the Directive on Automated Decision-Making and the Algorithmic Impact Assessment (AIA) [5, 19, 114], the AI register reveals substantial variation in how departments interpret, operationalize, and document these requirements. Many entries emphasize efficiency, workload reduction, and technical capability, while providing limited detail on uncertainty, human oversight, or evaluative practices. Despite requirements to complete and publish AIA results [90], a study found that only 4% of 303 automated tools had publicly available government assessments [4]. Our study exposes the persistent gap between formal obligations and visibility in practice: even when systems are disclosed, they often omit governance-critical information, including degrees of uncertainty, training arrangements, and discretion configuration, while departmental latitude to reinterpret central directives produces heterogeneous reporting

standards. Viewed against stalled efforts to regulate private-sector AI through binding legislation such as the proposed AI and Data Act [7, 62, 129], the GC AI Register appears as a compensatory and not an exemplary governance artifact.

Simultaneously, the register exposes the infrastructural dependencies that complicate and highlight the persistent challenges associated with Canada's claims of AI compute sovereignty. A substantial proportion of public-sector systems relies on tools developed by transnational vendors, notably Microsoft, revealing how AI deployment is embedded in the political landscape, technical expertise, global flows of capital, data, labor, compute infrastructure, and natural resources [8]. Reliance on externally developed tools, as revealed by the GC's own AI repository, indicates that Canada's pursuit of sovereign AI is structurally constrained beyond domestic policy, undermining claims of AI sovereignty and technological autonomy [18]. These tensions are further illustrated by partnerships emerging from Canada's sovereign compute investments, where federally funded domestic firms collaborate with foreign entities to develop a "fully sovereign AI factory" using NVIDIA graphics processing units (GPUs) and Hewlett Packard Enterprise computing infrastructure [65, 104]. Overall, our study highlights the breadth of infrastructural, organizational, and transnational considerations the GC must account for in shaping its AI strategies and ambitions.

5.2 Implication for Algorithmic Accountability Research

Based on our findings, we propose that algorithmic accountability initiatives should be understood as instruments of **ontological design**. It means that sociomaterial practices not only create objects or systems but also actively shape the kinds of beings, relationships, and ways of knowing that can exist in the world [137]. From this view, the GC AI register does not merely describe algorithmic systems but actively shapes which systems, decisions, and responsibilities are recognized as governable. Our analysis suggests that the register enacts a particular ontology of AI systems—one that foregrounds generative AI and LLMs while rendering other forms of AI peripheral; an ontology of AI objectives centered on efficiency and workload reduction; and an ontology of algorithmic accountability that nominally assigns responsibility to human decision-makers while obscuring how algorithmic automation conditions those decisions. Such disclosure practices ontologically impact how definitional boundaries delimit accountability. Viewed in conjunction with the ADMAPS framework, the unarticulated concerns of discretion, expertise, and uncertainty do not disappear; instead, they are shifted away from algorithmic accountability and absorbed by frontline practitioners.

These ontological effects are most visible at the boundaries of accountability artifacts, such as the GC AI register. For example, the absence of IRCC's Chinook from the register, despite its documented role in processing temporary resident visas, study and work permit applications [77, 133], illustrates how definitional and classificatory choices delimit what becomes accountable. Despite widely raised concerns about Chinook [23, 66], it does not appear as a named system in the Register. Whether intentionally excluded or described under a different name, this non-appearance is consequential: it renders a materially impactful form of algorithmic

mediation illegible to formal governance mechanisms. From an ontological design perspective, this is not merely a gap in coverage but an active configuration of accountability. By failing to account for how AI systems are known, named, and referenced by stakeholders outside government agencies, the GC AI Register reinforces a narrow ontology of algorithmic accountability. Taken together, such disclosure infrastructures do not simply fall short of comprehensive oversight; they actively shape and are shaped by the overlap of vocabularies and terminologies through which different stakeholders recognize, contest, and govern accountability.

5.3 Implications for Public Trust and Governance

Beyond Canadian AI policy and algorithmic accountability research, our findings raise broader concerns about public trust and democratic governance. Public-sector AI systems increasingly operate in high-stakes domains, such as immigration, social services, and security [21, 39, 67], where legitimacy depends not only on procedural compliance but on citizens' ability to understand, contest, and seek redress for algorithmically mediated decisions, and hence, transparency mechanisms are central to sustaining public trust [132]. Our analysis complicates this assumption by showing that partial, stylized, or selective disclosure can produce the appearance of accountability without substantive guarantees. For affected individuals, such representations offer little insight into how decisions are made, where discretion resides, or how harms might be challenged. Such disclosure practices risk shifting accountability onto the public without providing the visibility or institutional means needed to exercise it, thereby weakening the democratic promise of transparency [2, 96]. These practices may normalize procedural or symbolic accountability and bolster existing research's concern about the obfuscation of data-driven decision-making tools [110]. When registers function as endpoints rather than entry points for scrutiny, they may stabilize public trust at the surface level. However, in public institutions, where algorithmic systems directly shape access to rights, services, and protections, such arrangements risk eroding trust over time. Taken together, we argue that the public value of governance artifacts, such as the GC AI register, lies not in their mere availability but in their capacity to enable recognition, contestation, and accountability.

6 Conclusion

In this paper, we analyzed the complete Canadian federal Public AI Register through the lens of ADMAPS, examining it as a governance artifact rather than a neutral transparency mechanism. We showed how the register produces specific bureaucratic silences, especially around discretion, infrastructural dependence, and uncertainty, and we argued that AI registers function as instruments of ontological design. On this basis, we outlined design implications for future register schemes that aim to support democratic oversight rather than only creating the appearance of transparency. Instead of viewing such artifacts as endpoints of transparency, stakeholders (e.g., policymakers and civil society actors) should engage with these as sites of ongoing accountability. For algorithmic accountability and transparency researchers, this means moving beyond evaluating whether disclosure exists or how much information it provides to

sociotechnically examine how governance artifacts actively configure who and what can be held accountable.

7 Endmatter

7.1 Generative AI Usage Disclosure Statement

We used the writing-assistant software, Grammarly Premium (which includes OpenAI and Azure as part of its technical stack [47]), to improve clarity and sentence structure in our paper.

7.2 Author Positionality

Prior scholarship in social computing, FAcCT, computational linguistics, and other adjacent fields has highlighted how researchers' positionality shapes their motivations and perspectives [9, 15, 29]. The first author is a Bangladeshi immigrant in North America and cisgender man. The second author is a Haitian-Canadian and a cisgender woman. The third author is a Bangladeshi-Canadian and cisgender man. The fourth author is an Indian-Canadian and cisgender man. We come from interdisciplinary backgrounds, including computer science, economics, information science, and political science. Our decision to examine accountability in the Canadian Government's AI register stems from our positionalities in terms of our research interests, backgrounds, and concentration in algorithmic fairness, tech policies, human-centered data science, and critical human-computer interaction.

Acknowledgments

This study was partially supported by the Institute of Health Emergencies and Pandemics postdoctoral fellowship, the School of Cities Urban Challenge Grant at the University of Toronto, and an NSERC Discovery Grant.

References

- Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *new media & society* 20, 3 (2018), 973–989.
- Cecilia Aragon, Shion Guha, Marina Kogan, Michael Muller, and Gina Neff. 2022. *Human-centered data science: an introduction*. MIT Press.
- Starling Centre at Western University. 1 Fe. Tracking Automated Government 'TAG' Register Canada. <http://tagregistercanada.ca/>
- Blair Attard-Frost. 2022. Comments on the 3rd Review of the Directive on Automated Decision-Making. Available at SSRN 4238164 (2022).
- Blair Attard-Frost, Ana Brandusescu, and Kelly Lyons. 2024. The governance of artificial intelligence in Canada: Findings and opportunities from a review of 84 AI governance initiatives. *Government Information Quarterly* 41, 2 (June 2024), 101929. doi:10.1016/j.giq.2024.101929
- Blair Attard-Frost, Ana Brandusescu, David Gray Widder, and Christelle Tessono. 2025. AI Countergovernance: Lessons Learned from Canada and Paris | Tech-Policy.Press. <https://techpolicy.press/ai-countergovernance-lessons-learned-from-canada-and-paris>
- Blair Attard-Frost and David Gray Widder. 2025. The ethics of AI value chains. *Big Data & Society* 12, 2 (June 2025), 20539517251340603. doi:10.1177/20539517251340603
- Mariam Attia and Julian Edge. 2017. Be (com) ing a reflexive researcher: a developmental approach to research methodology. *Open Review of Educational Research* 4, 1 (2017), 33–45.
- Stefan Baack. 2024. A critical analysis of the largest source for generative ai training data: Common crawl. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 2199–2208.
- Ricardo Baeza-Yates. 2020. Bias in search and recommender systems. In *Proceedings of the 14th ACM Conference on Recommender Systems*. 2–2.
- Nagadivya Balasubramaniam, Marjo Kauppinen, Antti Rannisto, Kari Hiekkänen, and Sari Kujala. 2023. Transparency and explainability of AI systems: From ethical guidelines to requirements. *Information and Software Technology* 159 (2023), 107197.
- Yoshua Bengio, Sören Mindermann, Daniel Privitera, Tamay Besiroglu, Rishi Bommasani, Stephen Casper, Yejin Choi, Philip Fox, Ben Garfinkel, Danielle Goldfarb, et al. 2025. International AI safety report. *arXiv preprint arXiv:2501.17805* (2025).
- Reuben Binns. 2018. Algorithmic accountability and public reason. *Philosophy & technology* 31, 4 (2018), 543–556.
- Steven Bird. 2020. Decolonising speech and language technology. In *Proceedings of the 28th international conference on computational linguistics*. 3504–3519.
- William Boag, Harini Suresh, Bianca Lepe, and Catherine D'Ignazio. 2022. Tech worker organizing for power and accountability. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 452–463.
- Ana Brandusescu. 2021. *Artificial Intelligence Policy and Funding in Canada: Public Investments, Private Interests*. Technical Report. Centre for Interdisciplinary Research on Montreal, McGill University.
- Ana Brandusescu. 2025. Transjurisdictional AI Governance and the Myth of Sovereign AI. Invited presentation at the Taiwan–Canada Frontiers of Science Symposium. Co-hosted by CIFAR, DSET, and NSTC.
- Ana Brandusescu and Renee Sieber. 2022. Comments on the third review of Canada's Directive on Automated Decision-Making. Available at SSRN 4271913 (2022).
- Ana Brandusescu and Renée E. Sieber. 2025. Missed opportunities in AI regulation: lessons from Canada's AI and data act. *Data & Policy* 7 (Jan. 2025), e40. doi:10.1017/dap.2025.17
- Dennis Broeders and Huub Dijstelbloem. 2015. The Datafication of mobility and migration management: The mediating state and its consequences. In *Digitizing identities*. Routledge, 242–260.
- Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward algorithmic accountability in public services: A qualitative study of affected community perspectives on algorithmic decision-making in child welfare services. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- Lisa Ruth Brunner and Wei William Tao. 2024. Artificial intelligence and automation in the migration governance of international students: An accidental ethnography. *Journal of International Students* 14, 1 (2024), 269–288.
- Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- Stephen Casper, Carson Ezell, Charlotte Siegmann, Noam Kolt, Taylor Lynn Curtis, Benjamin Bucknall, Andreas Haupt, Kevin Wei, Jérémy Scheurer, Marius Hobbhahn, et al. 2024. Black-box access is insufficient for rigorous AI audits. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 2254–2272.
- Jennifer Cobbe, Michael Veale, and Jatinder Singh. 2023. Understanding accountability in algorithmic supply chains. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1186–1197.
- Nicholas Kluge Corrêa, Camila Galvão, James William Santos, Carolina Del Pino, Edson Pontes Pinto, Camila Barbosa, Diogo Massmann, Rodrigo Mambri, Luiza Galvão, Edmund Terem, et al. 2023. Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. *Patterns* 4, 10 (2023).
- Dipto Das, Shion Guha, Jed R Brubaker, and Bryan Semaan. 2024. The "COLonial Impulse" of Natural Language Processing: An Audit of Bengali Sentiment Analysis Tools and Their Identity-based Biases. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–18.
- Dipto Das, Shion Guha, and Bryan Semaan. 2023. Toward cultural bias evaluation datasets: The case of Bengali gender, religious, and national identity. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*. 68–83.
- Dipto Das, Carsten Østerlund, and Bryan Semaan. 2021. "Jol" or "Pani"? How Does Governance Shape a Platform's Identity? *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–25.
- Lina Dencik, Arne Hintz, Joanna Redden, and Harry Warne. 2018. Data scores as Governance: Investigating uses of citizen scoring in public services project report. (2018).
- Wesley Hanwen Deng, Boyuan Guo, Alicia Devrio, Hong Shen, Motahhare Eslami, and Kenneth Holstein. 2023. Understanding practices, challenges, and opportunities for user-engaged algorithm auditing in industry practice. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18.
- Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, and Hilary Nicole. 2021. On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society* 8, 2 (2021), 20539517211035955.
- Daria Dergacheva, Vasilisa Kuznetsova, Rebecca Scharlach, and Christian Katzenbach. 2023. One Day in Content Moderation: Analyzing 24 h of Social Media Platforms' Content Decisions through the DSA Transparency Database.

- (2023).
- [35] Nicholas Diakopoulos. 2015. Algorithmic accountability: Journalistic investigation of computational power structures. *Digital journalism* 3, 3 (2015), 398–415.
- [36] Nicholas Diakopoulos. 2016. Accountability in algorithmic decision making. *Commun. ACM* 59, 2 (2016), 56–62.
- [37] Nicholas Diakopoulos. 2025. Prospective Algorithmic Accountability and the Role of the News Media. In *Computer Ethics Across Disciplines: Deborah G. Johnson and Algorithmic Accountability*. Springer, 125–143.
- [38] Chiara Patricia Drolsbach and Nicolas Pröllochs. 2024. Content moderation on social media in the EU: Insights from the DSA Transparency Database. In *Companion Proceedings of the ACM Web Conference 2024*. 939–942.
- [39] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. Macmillan+ ORM.
- [40] Florian Eyert and Paola Lopez. 2023. Rethinking transparency as a communicative constellation. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 444–454.
- [41] Norman Fairclough. 2023. Critical discourse analysis. In *The Routledge handbook of discourse analysis*. Routledge, 11–22.
- [42] Angela Fan and Claire Gardent. 2022. Generating Biographies on Wikipedia: The Impact of Gender Bias on the Retrieval-Based Generation of Women Biographies. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 8561–8576. doi:10.18653/v1/2022.acl-long.586
- [43] Heike Felzmann, Eduard Fosch Villaronga, Christoph Lutz, and Aurelia Tamò-Larrieux. 2019. Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. *Big Data & Society* 6, 1 (2019), 2053951719860542.
- [44] Ben Gansky and Sean McDonald. 2022. CounterFAcCTual: How FAcCT undermines its organizing principles. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1982–1992.
- [45] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé II, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (2021), 86–92.
- [46] Mélanie Gornet, Simon Delarue, Maria Boritchev, and Tiphaine Viard. 2024. Mapping AI ethics: a meso-scale analysis of its charters and manifestos. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 127–140.
- [47] Grammarly. n.d.. Does Grammarly use Sub-processors? <https://support.grammarly.com/hc/en-us/articles/360036884632-Does-Grammarly-use-Sub-processors>. [Accessed: Jan 10, 2026].
- [48] Samuel Groesch, Alena Birrer, Natascha Just, and Florian Saurwein. 2025. Big data, small answers: How the DSA Transparency Database falls short of its regulatory objectives. *Telecommunications Policy* (2025), 103088.
- [49] M Haataja, L van de Fliert, and P Rautio. 2020. Public AI Registers: Realising AI transparency and civic participation in government use of AI. *Whitepaper Version 1* (2020).
- [50] Thilo Hagendorff. 2020. The ethics of AI ethics: An evaluation of guidelines. *Minds and machines* 30, 1 (2020), 99–120.
- [51] Charlotte Högberg. 2024. Stabilizing translucencies: Governing AI transparency by standardization. *Big Data & Society* 11, 1 (2024), 20539517241234298.
- [52] Mohammad Hossein Jarrahi, Ali Memariani, and Shion Guha. 2023. The Principles of Data-Centric AI. *Commun. ACM* 66, 8 (July 2023), 84–92. doi:10.1145/3571724
- [53] Deborah G Johnson. 2021. Algorithmic accountability in the making. *Social Philosophy and Policy* 38, 2 (2021), 111–127.
- [54] Nari Johnson and Hoda Heidari. 2023. Assessing AI Impact Assessments: A Classroom Study. *arXiv preprint arXiv:2311.11193* (2023).
- [55] Rishabh Kaushal, Jacob Van De Kerkhof, Catalina Goanta, Gerasimos Spanakis, and Adriana Iamnitchi. 2024. Automated transparency: A legal and empirical analysis of the digital services act transparency database. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1121–1132.
- [56] Pat Kelly. 2022. *Facial recognition technology and the growing power of artificial intelligence*. Parliament.
- [57] Svetlana Kiritchenko and Saif Mohammad. 2018. Examining Gender and Race Bias in Two Hundred Sentiment Analysis Systems. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*. Association for Computational Linguistics, New Orleans, Louisiana, 43–53. doi:10.18653/v1/S18-2005
- [58] Rob Kitchin and Gavin McArdle. 2016. What makes Big Data, Big Data? Exploring the ontological characteristics of 26 datasets. *Big data & society* 3, 1 (2016), 2053951716631130.
- [59] Katarzyna Klafkowska-Waśniowska and Jakub Drag. [n. d.]. In Search of Reason. Are Statements of Reason Under the Digital Services Act a Valuable Source of Information on Content Moderation Practices? Available at SSRN 5257655 ([n. d.]).
- [60] Joshua A Kroll. 2021. Outlining traceability: A principle for operationalizing accountability in computing systems. In *Proceedings of the 2021 ACM Conference on fairness, accountability, and transparency*. 758–771.
- [61] Blaine Kuehnert, Rachel Kim, Jodi Forlizzi, and Hoda Heidari. 2025. The “Who,” “What,” and “How” of Responsible AI Governance: A Systematic Review and Meta-Analysis of (Actor, Stage)-Specific Tools. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 2991–3005.
- [62] Mahtab Laghaei, Christelle Tessono, Viet Vu, and André Côté. 2025. *Submission to the Consultation on Canada’s Renewed AI Strategy*. Technical Report. The Dais, University of Toronto. <https://dais.ca/reports/submission-to-the-consultation-on-canadas-renewed-ai-strategy/>
- [63] Stefan Larsson and Fredrik Heintz. 2020. Transparency in artificial intelligence. *Internet policy review* 9, 2 (2020), 1–16.
- [64] Staffan I Lindberg. 2013. Mapping accountability: core concept and subtypes. *International review of administrative sciences* 79, 2 (2013), 202–226.
- [65] Tricia Lo. 2025. TELUS opens Canada’s first fully Sovereign AI Factory. <https://www.telus.com/en/about/news-and-events/media-releases/telus-opens-canadas-first-fully-sovereign-ai-factory>
- [66] Shaina Luck. 2025. Immigration lawyers concerned IRCC’s use of processing technology leading to unfair visa refusals. <https://www.cbc.ca/news/canada/nova-scotia/immigration-canada-ircc-technology-1.7632130>. [Accessed: Dec 28, 2025].
- [67] Kristian Lum and William Isaac. 2016. To predict and serve? *Significance* 13, 5 (2016), 14–19.
- [68] James G March. 1978. Bounded rationality, ambiguity, and the engineering of choice. *The bell journal of economics* (1978), 587–608.
- [69] Tasmiah Tahsin Mayeesha, Farzana Islam, and Nova Ahmed. 2024. AI4Bangladesh: AI Ethics for Bangladesh—Challenges, Risks, Principles, and Suggestions. In *Proceedings of the 13th International Conference on Information & Communication Technologies and Development*. 260–272.
- [70] Kelly McConvey and Shion Guha. 2024. “This is not a data problem”: Algorithms and Power in Public Higher Education in Canada. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [71] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–23.
- [72] Angelina McMillan-Major, Emily M Bender, and Batya Friedman. 2024. Data statements: From technical concept to community practice. *ACM Journal on Responsible Computing* 1, 1 (2024), 1–17.
- [73] Danaë Metaxa, Joon Sung Park, Ronald E Robertson, Karrie Karahalios, Christo Wilson, Jeff Hancock, Christian Sandvig, et al. 2021. Auditing algorithms: Understanding algorithmic systems from the outside in. *Foundations and Trends® in Human-Computer Interaction* 14, 4 (2021), 272–344.
- [74] Michael Meyer and Ruth Wodak. 2001. *Methods of critical discourse analysis*. Sage London.
- [75] Microsoft. n.d.. Canada Protected B. <https://learn.microsoft.com/en-us/azure/databricks/security/privacy/cccs-medium-protected-b>. [Accessed: December 13, 2025].
- [76] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. 220–229.
- [77] Petra Molnár. 2024. *The Walls Have Eyes: Surviving Migration in the Age of Artificial Intelligence* (first edition. ed.). New Press, The, New York.
- [78] Erina Seh-Young Moon, Devansh Saxena, Dipto Das, and Shion Guha. 2025. The Datafication of Care in Public Homelessness Services. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [79] Maya Murad. 2021. *Beyond the “Black Box”: Enabling Meaningful Transparency of Algorithmic Decision-Making Systems through Public Registers*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [80] Esther Nieuwenhuizen. 2024. Algorithm Registers: A Box-Ticking Exercise or Meaningful Tool for Transparency? *Information Polity* 29, 4 (2024), 415–433.
- [81] Chris Norval, Kristin Cornelius, Jennifer Cobbe, and Jatinder Singh. 2022. Disclosure by Design: Designing information disclosures to support meaningful transparency and accountability. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 679–690.
- [82] Government of Bangladesh. 2024. National Artificial Intelligence Policy. <https://tinyurl.com/BDAIPolicy>. Accessed: December 21, 2025.
- [83] Government of Canada. [n. d.]. Canadian Sovereign AI Compute Strategy. <https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy>
- [84] Government of Canada. 2000. Personal Information Protection and Electronic Documents Act (PIPEDA). <https://laws-lois.justice.gc.ca/eng/acts/P-8.6/>. (2000). [Accessed: March 12, 2026].
- [85] Government of Canada. 2021. Directive on Automated Decision-Making. <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592>. [Accessed: March 15, 2026].
- [86] Government of Canada. 2024. Pan-Canadian Artificial Intelligence Strategy. <https://ised-isde.canada.ca/site/ai-strategy/en/pan-canadian-artificial-intelligence-strategy> Last Modified: 2024-12-03 Publisher: Innovation, Science and Economic Development Canada.

- [87] Government of Canada. 2025. Departments and agencies. <https://www.canada.ca/en/government/dept.html>. [Accessed: December 12, 2025].
- [88] Government of Canada. 2025. Government of Canada AI Register (Minimum Viable Product) - GC AI Register (MVP). <https://open.canada.ca/data/en/dataset/fcbc0200-79ba-4fa4-94a6-00e32face6b/resource/369f6f34-148a-42ed-b581-8c164e941a89>. [Accessed: December 16, 2025].
- [89] Government of Canada. 2025. What is a Labour Market Impact Assessment? <https://ircc.canada.ca/english/helpcentre/answer.asp?qnum=163&top=17>. [Accessed: Jan 1, 2026].
- [90] Treasury Board Secretariat of Canada. 2024. Directive on Automated Decision-Making. <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592> Last Modified: 2023-04-25.
- [91] Government of Canada policy. 2026. Algorithmic Impact Assessment tool. <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html>. [Accessed: March 15, 2026].
- [92] Victor Ojewale, Harini Suresh, and Suresh Venkatasubramanian. 2026. Audit Trails for Accountability in Large Language Models. *arXiv preprint arXiv:2601.20727* (2026).
- [93] Chinasa T Okolo. 2023. AI in the Global South: Opportunities and challenges towards more inclusive governance. *The Brookings Institution* (2023).
- [94] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kicimcan. 2019. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in big data* 2 (2019), 13.
- [95] Global Partnership on Artificial Intelligence. 2024. Algorithmic Transparency in the Public Sector: A state-of-the-art report of algorithmic transparency instruments. <https://wp.oecd.ai/app/uploads/2025/05/algorithmic-transparency-in-the-public-sector.pdf>. [Accessed: Dec 28, 2025].
- [96] Frank Pasquale. 2015. The black box society: The secret algorithms that control money and information. In *The black box society*. Harvard university press.
- [97] Lindsay Poirier. 2021. Reading datasets: Strategies for interpreting the politics of data signification. *Big Data & Society* 8, 2 (2021), 20539517211029322.
- [98] Lindsay Poirier. 2022. Accountable data: The politics and pragmatics of disclosure datasets. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1446–1456.
- [99] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. 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.
- [100] Joanna Redden. 2018. Democratic governance in an age of datafication: Lessons from mapping government discourses and practices. *Big Data & Society* 5, 2 (2018), 2053951718809145.
- [101] Joanna Redden, Jessica Brand, Ina Sander, and Harry Warne. 2022. Datafied child welfare services as sites of struggle. *Data justice and the right to the City* (2022), 69.
- [102] Joanna Redden, Jessica Brand, Ina Sander, Harry Warne, A Grant, and D White. 2022. Automating public services: Learning from cancelled systems. *Cardiff, UK: Data Justice Lab, Cardiff University* 30 (2022).
- [103] Joanna Redden, Lina Dencik, and Harry Warne. 2020. Datafied child welfare services: unpacking politics, economics and power. *Policy Studies* 41, 5 (2020), 507–526.
- [104] Alex Riehl. 2025. Telus partners with Nvidia to transform its data centre into a Canadian "Sovereign AI Factory". <https://betakit.com/telus-partners-with-nvidia-to-launch-canadian-sovereign-ai-factory/>
- [105] Jonathan Roberge, Guillaume Dandurand, Kevin Morin, and Marius Senneville. 2022. Les narvals et les licornes se cachent-ils pour mourir ? De la cybernétique, de la gouvernance et d'Element AI. *Réseaux* 232233, 2 (June 2022), 169–196. doi:10.3917/res.232.0169 Publisher: La Découverte.
- [106] Ronald E Robertson, Shan Jiang, Kenneth Joseph, Lisa Friedland, David Lazer, and Christo Wilson. 2018. Auditing partisan audience bias within google search. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018).
- [107] Christian Sandvig, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry* 22, 2014 (2014), 4349–4357.
- [108] Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A Smith. 2019. The risk of racial bias in hate speech detection. In *Proceedings of the 57th annual meeting of the association for computational linguistics*. 1668–1678.
- [109] Beatrice Savoldi, Marco Gaido, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2022. Under the Morphosyntactic Lens: A Multifaceted Evaluation of Gender Bias in Speech Translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 1807–1824. doi:10.18653/v1/2022.acl-long.127
- [110] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2020. A human-centered review of algorithms used within the US child welfare system. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [111] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2021. A framework of high-stakes algorithmic decision-making for the public sector developed through a case study of child-welfare. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–41.
- [112] Devansh Saxena and Shion Guha. 2024. Algorithmic harms in child welfare: Uncertainties in practice, organization, and street-level decision-making. *ACM Journal on Responsible Computing* 1, 1 (2024), 1–32.
- [113] Devansh Saxena, Erina Seh-Young Moon, Aryan Chaurasia, Yixin Guan, and Shion Guha. 2023. Rethinking "Risk" in Algorithmic Systems Through A Computational Narrative Analysis of Casenotes in Child-Welfare. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–19.
- [114] Teresa Scassa. 2021. Administrative Law and the Governance of Automated Decision-Making: A Critical Look at Canada's Directive on Automated Decision-Making. *UBC Law Review* 54, 1 (2021). doi:10.2139/ssrn.3722192
- [115] Morgan Klaus Scheuerman, Allison Woodruff, and Jed R Brubaker. 2025. How Data Workers Shape Datasets: The Role of Positionality in Data Collection and Annotation for Computer Vision. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (2025), 1–42.
- [116] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*. 59–68.
- [117] Shilad Sen, Margaret E Giesel, Rebecca Gold, Benjamin Hillmann, Matt Lesicko, Samuel Naden, Jesse Russell, Zixiao Wang, and Brent Hecht. 2015. "Turkers, scholars," arafat" and "peace" cultural communities and algorithmic gold standards. In *Proceedings of the 18th acm conference on computer supported cooperative work & social computing*. 826–838.
- [118] Ushnish Sengupta. 2022. Algorithm registers as an infrastructural element for transparent government. *Artificial intelligence in the city: Building civic engagement and public trust* (2022), 68.
- [119] Hong Shen, Alicia DeVos, Motahhare Eslami, and Kenneth Holstein. 2021. Everyday algorithm auditing: Understanding the power of everyday users in surfacing harmful algorithmic behaviors. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–29.
- [120] Renée Sieber, Ana Brandusescu, and Jonathan van Geuns. 2026. Building AI Governance in Municipalities from the Ground Up.
- [121] Richard Simeon. 2006. The division of powers in Canada: an overview. (2006).
- [122] Jan Simson, Alessandro Fabris, and Christoph Kern. 2024. Lazy data practices harm fairness research. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 642–659.
- [123] Kacper Sokol and Peter Flach. 2020. Explainability fact sheets: A framework for systematic assessment of explainable approaches. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 56–67.
- [124] Ramya Srinivasan and Kanji Uchino. 2021. Biases in generative art: A causal look from the lens of art history. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 41–51.
- [125] Benjamin Lopez Steven. 2025. Ottawa launches AI register showing how the tech is being used within federal departments. <https://www.cbc.ca/news/politics/federal-government-public-ai-register-9.6995056>. [Accessed: December 21, 2025].
- [126] Lucy Suchman. 2002. Located accountabilities in technology production. *Scandinavian journal of information systems* 14, 2 (2002), 7.
- [127] Araz Taeihagh. 2025. Governance of generative AI. 22 pages.
- [128] Christelle Tessono, Nusrat Jahan Mim, and Shion Guha. 2025. Infrastructure and its Discontents: Contestations of Power in Canada's AI Governance Scaffold. (2025). Manuscript under review.
- [129] Christelle Tessono, Yuan Stevens, Momin M. Malik, Sonja Solomun, Supriya Dwivedi, and Sam Andrey. 2022. AI Oversight, Accountability and Protecting Human Rights. <https://dais.ca/reports/aida/>
- [130] Amaury Trujillo, Tiziano Fagni, and Stefano Cresci. 2025. The DSA Transparency Database: Auditing self-reported moderation actions by social media. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–28.
- [131] M Valderrama, MP Hermosilla, and R Garrido. 2023. State of the Evidence: Algorithmic Transparency. <https://www.opengovpartnership.org/wp-content/uploads/2023/05/State-of-the-Evidence-Algorithmic-Transparency.pdf>. *Open Government Partnership* (2023).
- [132] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–14.
- [133] Apurva Vohra. 2023. *Social order in the age of artificial intelligence: the use of technology in migration governance and decision-making*. Master's thesis. University of British Columbia. doi:10.14288/1.0437225
- [134] Algorithm Watch. 2025. AI Ethics Guidelines Global Inventory. <https://algorithmwatch.org/en/ai-ethics-guidelines-global-inventory/>. [Accessed: Dec 28, 2025].

- [135] Elizabeth Anne Watkins, Emanuel Moss, Jacob Metcalf, Ranjit Singh, and Madeleine Clare Elish. 2021. Governing algorithmic systems with impact assessments: Six observations. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 1010–1022.
- [136] Maranke Wieringa. 2020. What to account for when accounting for algorithms: a systematic literature review on algorithmic accountability. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 1–18.
- [137] Anne-Marie Willis. 2006. Ontological designing. *Design philosophy papers* 4, 2 (2006), 69–92.
- [138] Amy Winecoff and Miranda Bogen. 2025. Improving governance outcomes through AI documentation: Bridging theory and practice. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [139] Dora Zhao, Angelina Wang, and Olga Russakovsky. 2021. Understanding and evaluating racial biases in image captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14830–14840.
- [140] Anneke Zuiderwijk, Yu-Che Chen, and Fadi Salem. 2021. Implications of the use of artificial intelligence in public governance: A systematic literature review and a research agenda. *Government information quarterly* 38, 3 (2021), 101577.