

COOPÉRATION CANADA

In partnership with Canada

salanga

March 2026

Who Governs AI? A Trans Feminist Perspective on Who Gets a Say

Who Governs AI? A Trans Feminist Perspective on Who Gets a Say

Our third Equity & AI session invited participants into a critical examination of AI governance – not just what it is, but who shapes it, who participates in it, and who is left out. Facilitated by Blair Attard-Frost (she/her/they/them), Assistant Professor at the University of Alberta and Fellow at the Alberta Machine Intelligence Institute (Amii), the session explored how current AI governance systems are structured, why they often fail to prevent harm, and what becomes possible when we apply a trans feminist lens grounded in justice, inclusion, and agency.

Why This Session Matters

AI governance is often framed as a technical or policy challenge – something handled through regulations, standards, or ethical guidelines. But as this session made clear, governance is fundamentally about power. Who has the authority, resources, and legitimacy to shape AI systems? Whose knowledge counts? And who is expected to live with the consequences?

Current AI governance systems are not failing by accident. They are shaped by existing distributions of power – where state and industry actors dominate decision-making, and impacted communities are often excluded or consulted only superficially. Without addressing these underlying dynamics, governance efforts risk reinforcing the very harms they are intended to prevent.

What We Explored Together

Blair began by grounding the conversation in a shared understanding of AI and its lifecycle – from data collection to model development, deployment, and ongoing use. This framing made visible that AI is not a single tool, but a system shaped at multiple stages, each with governance implications. From there, the session moved quickly into an important reality: AI is already producing harm at scale. Because AI is now embedded across so many sectors, its impacts are not hypothetical. They are real, observable, and increasingly well-documented. Blair pointed to existing databases – including the AI Incident Database and the AIAAIC repository – which catalogue thousands of cases where AI systems have caused harm.

These incidents span a wide range of applications:

- Chatbots and generative AI tools
- Facial recognition systems
- Automated decision-making in hiring, finance, and policing
- Self-driving vehicles and robotics
- Education and healthcare technologies

Across these use cases, recurring concerns emerge – including bias and discrimination, safety failures, lack of transparency, and unreliable outputs. Importantly, these harms are not abstract. AI systems are contributing to physical, psychological, social, economic, and environmental impacts across different communities and contexts. This sets the foundation for the central question of the session: If harms are already widespread and well-documented, why are current AI governance systems so often failing to prevent them?

What a Trans Feminist Approach Reveals

To address this question, Blair introduced the analytical lens guiding her work: a trans feminist approach to AI governance. Rather than focusing only on technical performance or policy design, this lens draws attention to the ethical and relational dimensions of AI systems – particularly how power operates within them.

Trans and feminist theories of ethics emphasize principles such as justice, inclusivity, agency, community, and resistance. Applied to AI governance, this shifts the focus toward who is included in decision-making, whose knowledge is valued, and how systems can either reinforce or challenge existing inequalities. Blair's approach is informed by a range of thinkers and traditions, including:

- Work on trans community-based care, which centers mutual support, interdependence, and collective responsibility
- Glitch feminism, which challenges dominant norms embedded in digital systems and explores possibilities for refusal and disruption
- Community-led design and governance approaches, which prioritize participation, accountability, and lived experience in shaping technology

Together, this lens offers a different starting point for AI governance – one that begins not with systems, but with people, relationships, and power.

So, what do we mean by AI Governance?

Blair framed AI governance in broad, sociological terms – not just as policies or regulations, but as a system of practices intended to maximize benefits and minimize harms caused by AI systems. This definition is intentionally expansive. It recognizes that governance happens across many actors, decisions, and contexts – from technical design to policy to how systems are used and experienced in practice. Framing AI governance this way makes visible both its scope and complexity – and the many ways power and participation shape how these systems are governed.

One of the first insights Blair shared is that AI governance is often practiced through a set of recognizable organizational initiatives. Across the Canadian landscape, five common types of AI governance initiatives appear:

- Policies – statements of values and expectations that guide decision-making and resource allocation
- Programs – coordinated collections of projects with a shared strategic intent
- Strategic plans – detailed roadmaps outlining an organization's goals and priorities
- Standards – formal guidance developed through expert consensus on how systems should be designed, used, or managed
- Ethics statements – principles intended to guide responsible and ethical conduct

Together, these initiatives form much of what we typically recognize as "AI governance" in practice – shaping how organizations approach, manage, and justify their use of AI.

Key Takeaways - 1. AI governance in Canada is heavily shaped by industry priorities

Drawing from Blair's analysis of 84 AI governance initiatives across Canada, a clear pattern emerges: governance efforts tend to prioritize industry growth and technological advancement over broader social considerations. Most initiatives focus on:

- Industry and innovation – supporting the development and adoption of AI technologies
- Technology production and use – setting rules and allocating resources for how AI systems are built and deployed

In contrast, significantly less attention is given to:

- The impacts of AI on workers and the labour force
- Social systems and public services
- Education and training, including equipping people with the knowledge and skills to navigate AI safely and effectively

Blair highlighted that this reflects a broader dynamic: AI governance is often oriented toward enabling technological growth, while the social conditions needed to engage with, challenge, or benefit from these systems receive far less investment.

2. AI governance relies on unproven assumptions about "trickle-down" benefits

Blair highlighted that many AI governance initiatives are built on an underlying assumption: that supporting AI development and adoption in industry will ultimately lead to broad-based social benefit. This reflects a familiar "rising tide lifts all boats" logic – where innovation is expected to generate downstream benefits across society. However, as Blair emphasized, there is little clear evidence to support this assumption. In practice, this model relies on a series of speculative leaps:

- That technological innovation will translate into equitable economic gains
- That policy interventions will effectively steer outcomes toward public benefit
- That benefits will reach those most impacted by AI systems

Drawing on research into the political economy of AI, Blair noted that emerging evidence often points in the opposite direction – with AI adoption reinforcing existing inequalities and disproportionately impacting marginalized groups.

3. Public participation in AI governance is limited – and often exclusionary

Blair also pointed to a consistent pattern across governance initiatives: limited and uneven public participation. She explained that governance processes frequently involve targeted engagement with industry stakeholders, while broader public consultation is minimal or absent. As an example, Blair referenced Canada's proposed Artificial Intelligence and Data Act (AIDA). Introduced in 2022, the proposed legislation faced sustained criticism for its lack of meaningful public consultation – despite extensive closed-door engagement with industry actors. The bill ultimately failed to pass, illustrating deeper challenges in how governance processes are structured and who they include. Through these examples, Blair made visible a broader systemic issue: those most impacted by AI systems are often the least involved in shaping how they are governed.

4. Meaningful participation in AI governance requires resources – and access is unequal

Blair emphasized that meaningful participation in AI governance is resource-intensive. Engaging in governance processes requires:

- Technical and policy knowledge
- Time and sustained engagement
- Access to decision-making spaces
- Financial and organizational capacity

As a result, many groups – including community organizations, small unions, small businesses, and nonprofits – are often unable to participate in large-scale AI governance systems. Blair explained that this creates a structural imbalance: those with greater resources are better positioned to shape governance, while others are excluded from influencing decisions that directly affect them.

Even where public consultation exists, participation is not necessarily equitable. Marginalized groups often face barriers related to access, knowledge, and capacity – limiting their ability to engage meaningfully. This reveals a deeper constraint within current governance systems: participation is not just about being invited in – it is about having the resources and conditions needed to actually shape outcomes. **Across these patterns, Blair's analysis points to a deeper challenge: AI governance is not neutral – it reflects and reinforces existing distributions of power.**

What Can We Do About It?

To close the session, Blair invited participants to think differently about how AI governance itself is structured – contrasting top-down and bottom-up approaches. Most current AI governance systems follow a top-down model:

- Decision-making is centralized among governments and industry leaders
- Processes are large-scale and expert-driven
- Governance happens through formal mechanisms such as policies, laws, standards, and official guidance

While these approaches can provide structure and coordination, Blair emphasized that they often reinforce existing power dynamics – concentrating authority among those already resourced to participate. In contrast, Blair highlighted the potential of bottom-up AI governance:

- Distributed and decentralized – decision-making is shared across communities, practitioners, and diverse stakeholders
- Community-driven and smaller-scale – governance is shaped by those directly impacted, within specific contexts and workplaces
- More fluid and adaptive – including practices such as training, shared resources, community guidelines, and civic engagement (e.g., audits, petitions, public forums)

This is not about replacing one model entirely with another.

Rather, it is about recognizing that current systems are heavily weighted toward top-down approaches – and that more equitable governance requires greater investment in bottom-up, participatory models. Blair's broader call to action, grounded in her work on trans feminist AI governance, reinforces this shift:

- Expanding who is included in governance
- Challenging exclusionary systems
- Supporting community-led approaches that redistribute power

Together, these approaches point toward a different vision of AI governance – one that is not only more effective, but more accountable to the people most impacted by AI systems.

This session reinforces that AI governance is not just about creating better rules – it is about transforming who gets to shape those rules in the first place. If governance remains concentrated among state and industry actors, it will continue to reflect their priorities. Expanding participation – and resourcing it meaningfully – is essential to building governance systems that are not only more effective, but more just.

Key Resources

- Attard-Frost, B. & Lyons, K. (2024). AI governance systems: A multi-scale analysis framework, empirical findings, and future directions. AI and Ethics. https://link.springer.com/article/10.1007/s43681-024-00569-5
- From Attard-Frost, B. (2025). Transfeminist AI governance. First Monday, 30(4), April 2025. https://firstmonday.org/ojs/index.php/fm/article/view/14121
- Data Center Watch https://www.datacenterwatch.org/