November 19, 2026

A Day Late and a Dollar Short: Public Concern, Private AI, and State Obligations under the Human Right to Science

By Anna-Maria Hubert

Documents Commented On: *Canada's National Artificial Intelligence Strategy: AI for All* (Government of Canada, 2026)

> *"Many in the tech industry think they believe in democracy, but they don't think about it in terms of directly involving people in decisions. Partly because they think: 'People don't understand it,' and partly because they think 'This isn't possible,'" Divya Siddarth of Collective Intelligence Project told Democracy Technologies. "But every time we run something, we get the most interesting, engaged, thoughtful responses, and we show it can be done."*
>
> *— Graham Wetherall-Grujić, "Can Alignment Assemblies Bring Democracy to Silicon Valley?", Democracy Technologies (16 January 2024)*

Local opposition to data centres is arguably where citizens have had the greatest say on the development of Artificial Intelligence (AI) so far. Through zoning decisions, moratoria, litigation, and elections, communities have, in some cases, blocked or delayed projects worth billions. This opposition is driven by the material demands that data centres place on land, power, and water, but also by broader doubts about whether AI "is a force for good in society." However, siting a data centre is about as far downstream as decisions about AI go. By the time a community is asked to accept one, the questions that matter most — what the technology is for, whose interests it serves, and who bears its risks — have long since been determined elsewhere.

Formal efforts to involve the public in upstream decision-making about the development and governance of AI have been limited and recent, consisting mostly of one-off participatory and deliberative exercises, such as citizen juries (see here and here), citizen panels (see here), citizen assemblies (see here), and online consultations (see here), in a few jurisdictions, primarily led by governments, academics and civil society, and, on a few occasions, AI companies themselves (see here and here). Moreover, they have generally been intended to inform how AI may be regulated once it has been built, rather than to shape what the technology looks like in the first place. Those upstream decisions are made overwhelmingly within private companies (Stanford AI Index 2026). Although many in Silicon Valley may regard aligning AI with human interests as the defining challenge of our time, the question of whose interests, and what those interests are, has been left largely as an internal matter for the companies themselves. The public has been cast as a recipient of AI's benefits (and risks) rather than a participant in defining them. The central claim of this post is that this division of roles is not simply a policy failure, but a human rights problem.

Scholarship in science and technology studies (STS) has long stressed the public's role in shaping science and technology. This work has challenged the "deficit model," showing that citizens bring relevant knowledge and legitimate values to technical questions rather than an ignorance to be cured with better information (Wynne 1992). It has also called for public engagement to move "upstream," to the stage where research agendas are set and technological trajectories are still open (Wilsdon and Willis 2004; see also Stilgoe, Lock, and Wilsdon 2014). Jasanoff's "technologies of humility" spell out what such engagement should cover, including how a technology's objectives are framed; who is vulnerable and exposed to harm; whether risks and benefits are fairly distributed; and how institutions learn from their mistakes (Jasanoff 2003). The human right of everyone to enjoy the benefits of scientific progress and its applications, guaranteed principally in Article 15(1)(b) of the 1966 International Covenant on Economic, Social and Cultural Rights (ICESCR), gives these insights legal force. After lying fallow for decades, this "right to science" has attracted renewed attention. The UN Committee on Economic, Social and Cultural Rights (CESCR) elaborated its normative content in General Comment No. 25 (2020), and a growing body of scholarship has mapped its core elements: the right to participate in science and its applications, to access their benefits, and to be protected from their adverse effects. Jasanoff's concerns map closely onto these elements. My focus here is on the first, participation, and its relationship to the other two.

In this post, I argue that there is a fundamental gap between access to the benefits (and exposure to risks) of AI and citizens' ability to take part in the decisions that define those benefits and risks in the first place. Because the right to science is at its core participatory and anticipatory in nature, it requires states parties to inform the public, and to seek its input, on decisions about science and technology of consequence before they are widely deployed. The potential infringement of the right thus began well before the current debate over data centres. This duty falls primarily on the state, although private AI companies may also bear responsibilities. Ultimately, the problem is one of the legalization and institutionalization of the human right to science. The state's positive obligations to protect and fulfil the right require legislative, administrative, budgetary, and other measures that enable citizens to make their views known, and a framework that ensures those views are given effect. I conclude by considering what governments, and Canada in particular, must do to remedy the situation, with attention to Canada's new national AI strategy, *AI for All*.

The Scope of the Human Right to Science as Extending to AI Development

A threshold question is whether the human right to science extends beyond scientific research to technological innovation and development, which more aptly describes most current private AI activity. Article 15(1)(b) of the ICESCR protects the right of everyone to enjoy the benefits not only of scientific progress but also of "its applications." In General Comment No. 25, the CESCR confirmed that applications of science include technology deriving from scientific knowledge, such as information and communications technology (at para 7). The Committee's understanding of the scope of the right draws on the 2017 UNESCO Recommendation on Science and Scientific Researchers. The Recommendation treats research and development as a continuum, from fundamental and applied research through to "experimental development," and defines technology as knowledge relating directly to the production or improvement of goods or services (at para 1). The CESCR has further explained that the "benefits" of science refer first to the material results of the applications of scientific research (at para 8). Taken together, these sources arguably place

the development of AI systems within the scope of the right, especially given its current emergent and experimental nature.

The Right to Participate in Science and its Applications Under the Human Right to Science

> *"... science should be understood as a participatory good through and through, whether it is a matter of contributing to the scientific practice or enjoying its benefits later on."*
>
> *— Samantha Besson, "The Institutional Guarantee of the Human Right to Science" (2025) 25:1 HRLR at 8*

Under its first element, the human right to science includes a right to take part in decisions about science and its applications. The CESCR's General Comment No. 25 affirms that "the right of everyone to take part in cultural life includes the right of every person to take part in scientific progress and in decisions concerning its direction" (at para 10). It adds that "some decisions concerning the orientation of scientific research or the adoption of certain technical advancements should be subjected to public scrutiny and citizen participation," and that "[a]s far as possible, scientific or technological policies should be established through participatory and transparent processes and should be implemented with accompanying transparency and accountability mechanisms" (at para 55). The Committee thus makes clear that participation in the development and governance of science and technologies is part of the right itself. At least some decisions, including those about the orientation of research and the adoption of new technologies, must be opened to public scrutiny and input.

Yet the participatory element arguably remains the least developed part of the right, both in its normative content and in state practice. General Comment No. 25 prescribes no particular method of public engagement, for instance. It says only that methods giving the public an active role should be prioritized (at para 54). Nor have states, Canada included, done much to fill the implementation gap through legal and other measures, in accordance with their obligations under the human right to science, as argued further below.

This right to participate is also anticipatory. Decisions about "the adoption of certain technical advancements" are, by definition, made before those advancements are adopted (at para 55). General Comment No. 25 also ties participation directly to the control of risk and uncertainty. In controversial cases, the Committee observes, "the risks and potential of some technical advances … should be made public in order to enable society, through informed, transparent and participatory public deliberation, to decide whether or not the risks are acceptable" (at para 57). A right to help decide whether a technology's risks are acceptable can arguably only be exercised before that technology is widely deployed. By the time a community is asked to accept a data centre, that decision has long since been made. Participation confined to siting and resource concerns comes too late to satisfy the right.

The precautionary principle reinforces this anticipatory dimension. Under the principle, "in the absence of full scientific certainty, when an action or policy may lead to unacceptable harm to the public or the environment, actions will be taken to avoid or diminish that harm" (at para 56). The CESCR describes unacceptable harm as harm that is threatening to human life or health; serious

and irreversible; inequitable to present or future generations; or imposed without adequate consideration of the human rights of those affected (at para 56). Several of AI's potential harms arguably meet this threshold, from the water and energy demands of data centres to the more speculative, but potentially catastrophic or even existential risks that some developers themselves acknowledge. The final criterion is met where harms are imposed without any avenue for public participation in decisions about AI's development and deployment. On this definition, the exclusion of the public can itself render a harm unacceptable. The Committee is careful to add that the precautionary principle should not hinder scientific progress that benefits humanity (at para 57; see further Cass R Sunstein, "Beyond the Precautionary Principle" (2003) 151:3 *U Pa L Rev* 1003). Nevertheless, in controversial cases, of which AI is surely one, "participation and transparency become crucial" (at para 57).

States, AI Companies, and the Duty to Protect

> *"So often I'm in San Francisco talking to folks I know in the AI industry, and they'll say things, and in my head I'm just like, do you guys hear yourself? If any normal person walked into this conversation, they would flip out.*
>
> *One of those examples … is that question of, why would you build something if you think there's a 10 or 20 percent chance the thing you build kills us all? And I will hear people straight up say, yeah, that's true, but I think it's somewhat more likely, maybe there's a 50 percent chance, that we all live forever and all the diseases are cured. So the expected value of 50 percent we live forever outweighs the 20 percent we all die, and I'm willing to take that chance and roll the dice a few times.*
>
> *And I'm just like, man, I don't actually think that most people in this country or in the world would take that bet. But more than that, those folks don't really have a choice about whether they take that bet, because Silicon Valley is taking the bet and then imposing the costs and risks on other folks."*
>
> — *Jasmine Sun* in conversation with Nate Silver "*Why Does Everyone Hate Data Centers?", Silver Bulletin (17 August 2026)*

Like all human rights, the right to science imposes three types of obligation on states parties: to respect, to protect, and to fulfil (General Comment No. 25 at paras 41–50). The state is the primary duty-bearer. The obligation *to protect* requires states to prevent third parties, including private companies, from interfering with the enjoyment of the right, as affirmed by the UN Guiding Principles on Business and Human Rights (UNGP). In the context of AI, the state thus stands between the public and the companies developing the technology as the protector of their human rights.

States parties are not, as such, responsible for abuses committed by private AI companies. They are responsible where they fail to exercise due diligence, that is, where they fail to take appropriate steps to prevent, investigate, punish and redress such abuses (UNGP, Principle 1). AI companies, for their part, have their own responsibility *to respect* human rights (UNGP, Principle 11). It is unclear whether that responsibility extends to engaging the public directly in shaping the technology. The corporate responsibility to respect is framed largely in negative terms (John

Ruggie, Protect, Respect and Remedy: A Framework for Business and Human Rights, UN Doc A/HRC/8/5 (7 April 2008) at para 24), where companies must avoid causing or contributing to adverse human rights impacts through their own activities, and must carry out human rights due diligence to identify, prevent and mitigate those impacts and address them where they occur (UNGP, Principles 11, 13, and 17). Principle 18 does require that this due diligence "[i]nvolve meaningful consultation with potentially affected groups and other relevant stakeholders." But such consultation is a means of identifying risks to those affected, not necessarily a general right of the public to help shape a company's products. In the case of AI, however, the distinction may be of little importance in practice, since where potential impacts are as pervasive as AI's, the "potentially affected groups" may include almost everyone. Such engagement would, however, raise particular issues of legitimacy, accountability, and justice (e.g., through "participation washing" and issues of participatory design).

The primary human rights duties, including with regard to the participatory branch of the human right to science, remain on the state. In this vein, states' duty to protect also has a more specific implication for participation itself. States must safeguard public consultation and participation processes from undue influence by third parties (Romano and Boggio 2024 at 606). Where industry voices dominate the processes through which a technology's governance is decided, the duty to protect is engaged. I return to this point below in relation to Canada's recent consultation on AI transparency.

What This Means for the Realization of our Scientific Rights in Canada

Protection alone will not close the participation gap with respect to private AI development. What is required is the further legalization and institutionalization of the right to participate under the human right to science, including at the domestic level. This falls under the obligation *to fulfil*, the "most demanding of the three obligations" (Romano and Boggio 2024 at 606). At a minimum, this requires that states parties must give the right adequate recognition in their national political and legal systems, preferably through legislation. Some of this is not left to progressive realization. Amongst the core obligations that states parties must meet as a matter of priority, the CESCR's General Comment No. 25 lists the duty to "[d]evelop a participatory national framework law on [the right to science] that includes legal remedies in case of violations, and adopt and implement a participatory national strategy or action plan for the realization of this right" (at para 52). At present, Canada has neither.

The obligation to fulfil comprises duties to facilitate, provide, and promote the right (General Comment No. 12 at para 15; General Comment No. 14 at paras 33, 36–37; see also ICESCR, Article 15(2)). Each element of this tripartite structure offers a lens for assessing how the participatory branch of the right to science might be better realized within the context of AI development in Canada, including under Canada's new National Artificial Intelligence Strategy, *AI for All*.

Duty to facilitate. The duty *to facilitate* requires states parties to take positive measures that enable people to participate in AI development and governance, by creating the conditions, laws and institutions through which they can exercise the right. The problem is not that AI is failing to progress, indeed, it is advancing at breakneck speed. The problem is that it is advancing with little input from citizens into its design or governance. Facilitating participation would require laws,

institutions, and processes through which the public can meaningfully influence the direction of AI. Examples include standing citizens' assemblies or advisory bodies with a defined role in law and policy; participation requirements attached to public AI funding and procurement; and mandatory public consultation and impact assessment for major AI infrastructure.

A further question relates to the need to create processes through which public input can feed into the development of AI as well as its governance. In particular, this requires that the state not only ascertain the public's views, but also that it provides a pathway for its integration into the design and orientation, including presumably, since the realization of the right would require it, in privately developed AI. This already difficult task is further complicated by the fact that AI development is transnational, and thus more difficult to regulate by Canada on its own. This engages states parties' duty of international cooperation under Article 15(4) of the ICESCR. In this case, for example, the results of public engagement in Canada should inform Canada's positions in the international governance of AI.

AI for All does little of this. The government's one participatory exercise to date, a consultation on AI transparency that ran from July to 23 September 2026, dealt largely with transparency towards users of AI outputs. It did not ask the public how, or for what purposes, AI should be developed. As the Canadian Centre for Policy Alternatives (CCPA) notes, this narrow conception of transparency "stops at the screen" being "limited to how Canadians can know when they are interacting with AI systems and recognize AI generated content." Broader concerns "beyond the user experience", including job losses, privacy, misinformation and AI's environmental footprint, were largely left out.

The CCPA also warned that the process left "industry voices to shape a process meant to hold industry accountable," and that the government's discussion paper "pays little attention to how the AI industry itself resists transparency." Developers, for example, invoke the "black box" problem to claim that their models are too complex to explain, including what drives a given recommendation. These criticisms engage not only the duty to fulfil the right to participate, but also the duty, discussed above, to protect participatory processes from undue influence by third parties, including private companies.

Duty to provide. The duty *to provide* requires states parties to supply the means of participation where individuals or groups cannot participate on their own for reasons beyond their control. In practice, the state must ensure that people can take part on equal terms without discrimination. For example, meaningful participation depends on access to information that developers may otherwise withhold. States parties may therefore have to impose disclosure requirements on matters such as training data, model capabilities, and resource use. They may also have to provide funding and support for affected and marginalized communities, and accessible formats and languages, so that they can participate in public participation and engagement processes.

Here too the transparency consultation carried out by the federal government fell short. According to the CCPA, it left out precisely the information companies tend to shield as proprietary, such as the hidden labour, personal data and creative work used in training, and the water and energy use of data centres. Without that information, the public cannot participate fully. The survey was also criticized for being too technical for most Canadians. Those with policy or industry expertise thus potentially had greater influence over the public record, running counter to the duty to provide.

Duty to promote. The duty *to promote* requires states parties to raise awareness of the right to participate, and to build the critical and civic understanding needed to exercise it. Canada's *AI for All* does commit to a National AI Literacy Initiative providing entry-level AI training to all Canadians. But the initiative is framed around closing an adoption gap. In other words, it teaches people to use AI, but does not provide a pathway for them to have a say in what it looks like (beyond users' role in producing the data that is used to train and evaluate machine learning models). In this sense, the strategy is oriented towards the second branch of the human right to science, i.e., in promoting the sharing in the benefits of AI, rather than its first, which provides for a right of participation in scientific and technological progress, leaving the latter unfulfilled.

Concluding Thoughts

The core argument of this post relates to democratic rights in the development of science and technology, which increasingly affect people and the planet in profound and often disruptive ways. The participatory element of the right to science may be its least developed, both in law and in state practice. Yet its importance has never been greater, especially for high-stakes, uncertain, and controversial innovations such as AI. Indeed, participation is arguably a precondition for the right's other elements. People cannot meaningfully access the benefits of science, or be protected from its risks, unless those benefits and risks are defined transparently and with democratic input.

Implementing this element will take considerable work, particularly for private research, development and innovation. The private sector now performs roughly 70 percent of research and development in most OECD countries, much of it across borders, and the mechanisms for feeding democratic input into private activity remain ill-defined. The primary duty to build them lies with states, although businesses may bear responsibilities of their own.

In General Comment No. 25, the CESCR identifies core obligations that states parties must meet as a matter of priority. Amongst them, states must

> [d]evelop a participatory national framework law on [the right to science] that includes legal remedies in case of violations, and adopt and implement a participatory national strategy or action plan for the realization of this right that includes a strategy for the conservation, the development and the diffusion of science. (at para 52)

Canada, for its part, has neither. *AI for All* is a sectoral strategy for one technology, not a national plan for realizing the right to participate in science and its applications, and Canada has no framework law on the right at all. Here too the argument of this post extends well beyond AI. In general, Canada must adopt a participatory national framework law on the right to science, and identify the areas of scientific research and its applications where dedicated participatory mechanisms are most needed. AI, and the infrastructure it depends on, however, would be an obvious place to start.

This post may be cited as: Anna-Maria Hubert, "A Day Late and a Dollar Short: Public Concern, Private AI, and State Obligations under the Human Right to Science" (19

November 2026), online: ABlawg, http://ablawg.ca/wp-content/uploads/2026/11/Blog_AMH_AI_Right_To_Science.pdf

To subscribe to ABlawg by email or RSS feed, please go to http://ablawg.ca

Follow us on Twitter @ABlawg