

Tumbler Ridge – catalyst for AI regulation?

In the aftermath of the tragedy in Tumbler Ridge BC, when an 18-year old youth shot and killed eight people including five schoolchildren, and the revelation that the youth may have been interacting with ChatGPT to receive information, or even guidance, in regard to gun violence, the issue of safety of AI (artificial intelligence) systems has come sharply into focus.

OpenAI, the operator of ChatGPT, had identified through its internal procedures account activity by the 18-year old that it flagged as potentially unsafe. However, it did not report this information to law enforcement authorities because the relevant content did not meet its standard of an “imminent and credible risk of serious physical harm to others.”

Safety has been very much part of the public discourse on the future of AI and its role in advancing innovation. In 2022 the federal government tabled its Bill C-27 which, in addition to a new privacy law – the *Consumer Privacy Protection Act* – included its proposed *Artificial Intelligence and Data Act* (“AIDA”), requiring safety assessments of “high impact” AI systems. Those initiatives died on the order paper when Parliament was prorogued in January of last year.

Ethical AI regulation

The AI regulation regime proposed under AIDA followed closely the similar regime now in force in the EU - the EU’s *Artificial Intelligence Act* (“AI Act”), with one significant exception. While the EU regime seeks to regulate all “AI” systems, stipulating rules for those with risks ranging from minimal to high, AIDA focussed only on “high impact” systems having potential for significant risk. Both the EU Act and AIDA are what can be characterized as “horizontal” regulatory rules – i.e. they would apply across all industry sectors. This approach to regulation contrasts with a “sectoral” approach which contemplates industry-specific regulation targeted at identified AI actions that have potential for harm.¹

Like the EU AI Act, AIDA would have imposed obligations on developers and operators of AI systems to assess and mitigate potential risks of harm resulting from their use, and to the extent they cannot be mitigated, to cease providing a system. It also stipulated requirements for transparency and accountability for such systems, including posting publicly information regarding their policies for governance and risk mitigation and providing reports regarding their mitigation and safety plans to a designated regulatory authority.²

A [previous version of this article](#) was published in Law360 Canada, part of LexisNexis Canada Inc. March 9, 2026.

¹ See UK approach – delegation of rule-making to industry regulators - [A pro-innovation approach to AI regulation](#); UK Department of Science, Innovation and Technology; Feb. 2024.

² The Minister of Innovation, Science and Economic Development, acting through the Artificial Intelligence and Data Commissioner, who would be appointed by the Minister to carry out the Minister’s regulatory functions under AIDA.

The proposed AIDA regulatory framework focused on ensuring that systems operate safely and that any potential harms are mitigated before they occur. An important aspect of this safety thrust was the focus on public accountability requirements, which presumably should work to eliminate risks, both through an operator's internal due diligence procedures and by the scrutiny of such procedures by the regulator, and if necessary, the regulator's order to take steps to eliminate a risk including by ceasing to provide an AI service.

Current status of AI reform

With the failure of Bill C-27 to move forward before Parliament was prorogued, and with a new Liberal leadership in place, the federal government is now studying what approach should be taken going forward for AI oversight and, potentially, regulation.

Criticisms of AIDA, voiced both previously and currently, range from it being too heavy-handed and thus discouraging for innovation, to lacking sufficient consultation regarding the spectrum of potential harms from across civil society. Diverse approaches have been advanced, from taking a light touch to AI regulation focussing on sector-specific statutory rules adopted on a case-by-case basis,³ to a comprehensive regulatory framework addressing the evolving challenges of AI and digital platforms.

Commentators arguing for a pause on full-scope regulation have suggested that with on-going, rapid developments in AI technologies it is too early to anticipate a comprehensive regulatory scheme and, further, that undertaking a legislative adoption process at this time will introduce significant uncertainty at a pivotal moment in the development of such technologies.

To be noted, the focus on safety in the regulatory sphere has for the most part addressed issues of fairness, discrimination and security of information, as opposed to matters of physical safety. Notwithstanding this emphasis on psychological, not physical, risk considerations, initiatives to regulate AI evident to date, such as AIDA and the EU's AI Act, clearly include risk of physical harm as within scope. So, it may be asked, could AI regulation such as has been proposed under AIDA, have prevented the tragedy in Tumbler Ridge, had it been in force?

This question bears some analysis (see below). However the wider issue, now brought back into the current AI regulation debate, is whether broad-based regulation of AI such as AIDA would have represented, now should be advanced more actively than has been the case since the Bill C-27 version died (quietly) last year. One argument for advancing such omnibus-scope regulation, as opposed to more selective, sector-based regulation, would be that while certain AI risks such as discrimination and reputational harm can be identified and potentially addressed under existing regulatory regimes including employment and human rights, the innovative and fast-developing character of the technology will not be responsive to targeted, ex-ante regulatory initiatives seeking to anticipate and legislate for all potential risks in a particular field. At best, they likely would be responsive to ex-post experiences. Furthermore, while

³ See, for example Ontario [Working for Workers Four Act, 2024](#) requiring disclosure of AI application in assessing and hiring job applicants.

it may be possible to recognize many potential risks within defined areas of concern, such as human rights, employment and privacy, the potentially cross-cutting character of AI technologies belies the easy pigeon-holing of risks.

The Tumbler Ridge incident is a case in point. Prior to this incident, there likely was little consideration given to the potential use of general purpose AI systems such as ChatGPT by individuals considering a mass shooting. Or, if there might have been, at least by the operators of such systems (for example by their monitoring of the use, as appeared to be the case in this instance), there was no legislative public safety regime governing the ChatGPT AI system that could have laid down standards for risk identification and risk mitigation that the operator, OpenAI, would have been required to follow.

Online harms

An arguably sector-specific approach to regulating potential harms from AI may be found in the previous government's attempt to pass an *Online Harms Act*, directed primarily at social media platforms. A key focus of this proposed law, tabled in February 2024, was promoting online safety and reducing harms caused to individuals by harmful content online through social media, with a particular focus on children and youth. However it is not clear that the proposed law, as framed, would have extended to chatbots such as ChatGPT because the definition of social media in the bill required communication between two or more individuals. If the sector-specific approach to AI regulation ultimately wins out, we could see an online harms law that expressly extends to chatbots.⁴ Whether an online harms law or a more broadly applicable harms prevention law is the more appropriate approach to addressing such AI risks is a question for debate.

As with AIDA, the *Online Harms Act* would have imposed obligations on operators of regulated systems to assess and mitigate potential risks of harm, and to the extent the potential for harmful content cannot be mitigated, to block access.

The Tumbler Ridge incident is again instructive. The information is that the individual was interacting with the chatbot online and potentially learning about gun-related violence. It might be argued that such a development could have been addressed by an appropriately-wider scope online harms law – that would have clearly encompassed not only human-to-human online interactions but also machine-to-human interactions.

What risks?

However, prior to Tumbler Ridge, a sector-focussed AI risk-prevention regime addressing social media or online communications along the lines of the proposed *Online Harms Act*, would have included social/psychological issues affecting individuals but arguably not public/community risks extending to physical violence. It is clear that a significant limiting factor regarding the possible application of either of the proposed AI or online harms regimes to the Tumbler Ridge incident is that they were structured to protect individuals not the broader public, or communities at large.

⁴ See [Final Report to the AI Strategy Task Force](#), Taylor Owen; January 22, 2026.

Both proposed regimes focused on protecting against *harms affecting individuals*, whether as users of a system, or persons who may be affected negatively by the service provided by the system (such as an employment-related decision), or persons otherwise negatively impacted by the use of the service, such as by the posting of their non-consensual intimate images.

This focus can be seen in AIDA's description of the risks that it addresses - as potential *harms to an individual* - including physical or psychological harm, damage to property, and economic loss. One of the criticisms of the proposed AIDA regime was that it did not address potential harms to *groups of individuals*, or *community harms*. Now, in light of the Tumbler Ridge tragedy, a logical argument would be that any AI protective regime, going forward, should explicitly extend to potential public sphere risks.

In this regard, it may be noted that in a [November 28, 2023 letter](#) to the parliamentary Standing Committee on Industry and Technology (INDU) proposing an extension of the AIDA regime to general purpose (or generative) AI systems – not specifically addressed in the Bill as originally tabled – the then Minister of Innovation, Science and Industry recognized that the risks posed by these systems could extend to “societal scale risks”. However, the societal risks to which he was alluding were the risks posed by synthetic content potentially leading to disinformation and its attendant risks to democracy, not physical harms to members of the public.

Reporting requirement

To be noted is that neither proposed regime would have required an operator to report to a regulator or to law enforcement authorities any particular activity or use of a system, or content on the system, that it identified as potentially dangerous whether to the user or to third parties. In other words, had either of these proposed laws applied at the time prior to the Tumbler Ridge shootings, there would have been no *requirement* for OpenAI to report any use of its ChatGPT system that pointed to a potential shooting incident.

Therefore, consideration may be given to whether any AI/online content oversight regime should include a risk reporting requirement, whether to a regulator or to law enforcement. Such a rule could be problematic since it would imply extending mandatory surveillance to practically all private interactions involving a chatbot or social media. The alternative approach would be to focus on prevention – through enhanced public accountability and risk mitigation, building on these provisions in the previously proposed AI and online harms regulatory frameworks - which would entail advance identification of specific potential risks and, conceivably, reporting such risks to the appropriate authority.

For more information please contact: David Young 416-968-6286 david@dauidyounglaw.ca
Note: The foregoing does not constitute legal advice. © David Young