



February 9, 2024

Submission to the Standing Committee on Industry and
Technology on Bill C-27, *The Digital Charter
Implementation Act, 2022*

Alireza Kamyabi | Vancouver Coastal Health
Dr. Alex Choi | Vancouver Coastal Health
Kevin Lam | University of Toronto
Tina Purnat | University of Memphis
Juan Solorzano | Vancouver Coastal Health
Manik Saini | Vancouver Coastal Health
Dr. Geoffrey McKee | University of British Columbia
Dr. Ihoghosa Iyamu | University of British Columbia
Dr. Simon Carrol | University of Victoria
Dr. Lorie Donelle | University of South Carolina
Shannon Turner | Public Health Association of BC

Table of Contents

I. Summary	3
II. Potential Threats & Wellbeing Impacts of AI on Population Health	3
Embedding Systemic Bias	4
Exacerbating socioeconomic inequities.....	4
Eroding human-to-human interactions	5
Perpetuating Misinformation	5
III. List of Recommendations	6
IV. Recommendations to Committee	7
V. About PHABC	12
VI. Bibliography	13

I. Summary

This submission provides recommendations for revisions to the Artificial Intelligence and Data Act (AIDA), a pivotal piece of legislation under Bill C-27, which seeks to regulate the use, development, and deployment of AI systems in Canada. AIDA aims to strike a balance between fostering innovation and protecting consumers' privacy and rights in an increasingly digital and AI-driven landscape. However, from a population health perspective, there are gaps that need to be addressed to ensure the Act effectively safeguards population health and upholds principles of healthy equity.

The key concerns outlined in this document revolve around the definitions and scope of harm, biased outputs, and the regulatory reach of AIDA. The Act currently focuses on "high-impact systems," potentially overlooking the broader spectrum of AI applications that can directly, indirectly, or cumulatively affect the social determinants of health and consequently the health of our populations. Many experts point out that emerging AI systems, which are trained on current datasets and literature, often mirror existing biases in our society. If not addressed, this trend can continue and worsen the disparities already present, leading to poorer health outcomes, especially for disadvantaged groups (Norori, Hu, Aellen, & Tzovara, 2021). Moreover, a discernable discrepancy in access to digital technologies is emerging, amplifying the digital divide along a number of axes including age, ethnicity, socioeconomic status, and geographical location. While AI has the potential to enhance quality of life for Canadians, it is evident that realizing this potential will require regulatory action to mitigate harms and promote equitable outcomes across our communities.

Critically, the achievement of a system that works for all Canadians will require a diverse range of perspective in both development and regulation, particularly a focus on health equity. This inclusion is crucial to ensure that AI regulations are grounded in an understanding of the digital, commercial, and social determinants of health and the complex ways in which AI can impact health outcomes (Kickbusch & Holly, 2023).

In summary, while AIDA is a significant step towards regulating AI in Canada, it needs to have a broader scope with strengthened harm mitigation and accountability measures in place, among others, to protect the health of our communities.

II. Potential Threats & Wellbeing Impacts of AI on Population Health

Like the internet and social media, AI can be a double-edged sword, presenting both benefits and risks to population health and wellbeing. Implemented appropriately, AI could enhance population health responses through automation, extend our ability to reach underserved populations, deepen our insight into population health, and drive our ability to innovate. Already, AI applications are being used to conduct near-instantaneous literature reviews, develop new vaccines, and facilitate novel surveillance mechanisms. These advancements can boost healthcare capacity and increase accessibility while elevating standards of care, particularly in

the face of growing population health needs and limited resources. However, new AI tools also carry inherent risks, necessitating a thoughtful approach to mitigation. Given the profound societal shifts that AI can bring about, the negative repercussions of poorly managed AI policies can be monumental. It is crucial to effectively mitigate these adverse effects to fully harness AI's positive potential, acknowledging that achieving this balance is a complex yet essential task. As such, in this submission, consideration of the negative population health impacts is not meant to be a barrier to innovation. Rather, we argue that understanding the potential threats AI poses to population health will ensure that innovation truly serves the public interest, allowing us to devise strategies that might preserve the benefits of AI while effectively managing associated risks.

Embedding Systemic Bias

AI offers to support population health through enhanced data-driven approaches; yet, AI systems can also mask unchecked systemic biases that have the potential to exacerbate existing inequities. Bias in AI systems can stem from unrepresentative data samples during training (data-driven bias), human bias introduced during implementation (human bias), or inappropriate algorithmic weighting of different features within training datasets (algorithmic bias) (Norori, Hu, Aellen, & Tzovara, 2021).

Data-driven bias typically stems from the use of unrepresentative data samples during the initial stages of training. The exclusion of historically marginalized populations from data samples can result in AI systems demonstrating preferential accuracy towards privileged groups, reinforcing existing disparities, including among health outcomes. For instance, when AI-driven skin cancer detection algorithms are trained on datasets that are disproportionately comprise of patients with lighter skin tones, they demonstrate reduced accuracy for individuals with darker skin (Guo, Lee, Kassamali, Mita, & Nambudiri, 2022). Data-driven bias can also intensify societal prejudices if coupled with human biases when healthcare workers, influenced by conscious or unconscious racial biases, preferentially act on the algorithm's outputs (European Union Agency for Fundamental Rights, 2022).

Furthermore, AI systems are particularly vulnerable to bias due to their reliance on large, multi-faceted training datasets. Due to the sheer size of such datasets and the opaque nature of machine learning algorithms, these biases can often go unnoticed. Consequently, the explainability of AI systems, particularly in terms of the composition and characteristics of their training datasets, is a critical factor that must be emphasized when assessing the risks associated with their use. This involves ensuring transparency in how these datasets are curated and how they influence the behavior and decisions of AI systems.

Exacerbating socioeconomic inequities

Income is a pivotal determinant of health, with poverty increasing risks for mental illness, chronic disease, and lower life expectancy. AI's adoption is already showing worrying signs of exacerbating existing socioeconomic inequities by failing to compensate content creators whose

work has been used to train AI systems. Many current AI systems have been developed using the unpaid, unrecognized contributions of artists, authors, and labourers, transferring wealth to large corporations and owners of AI systems (Klein, 2023). Moreover, as AI is adopted in workplaces, privileged workers with the resources and training to rapidly adapt to and leverage AI have unfair advantages over those experiencing systemic barriers to adoption. This potential divide can further widen existing digital and socioeconomic disparities across age, education, and geographical location (e.g., rural, and remote). As such, AI-driven economic growth may disproportionately benefit privileged populations, thus exacerbating existing socioeconomic inequities and deepening wealth-driven health disparities. AI's environmental impacts (e.g., training a large language model can emit up to 300,000kg of CO₂ (Strubell, Ganesh, & McCallum, 2019)) can further amplify such inequities, given climate change disproportionately impacts low-income populations.

Moreover, the direction of AI research is largely influenced by major technology companies. Their decision-making, as seen in recent leadership challenges in OpenAI, often grapples with balancing commercial interests against broader public good. As such, it is critical to establish stringent accountability measures and incentive frameworks to ensure AI advancements are harnessed for the optimal benefit of population health.

Eroding human-to-human interactions

Social connections are an important determinant of health as they influence physical, mental, and emotional well-being (Hold-Lunstad, 2022). While technology has enabled us to communicate more effectively over time and space, the transition to digital, and often asynchronous, interactions via social media has impacted the quality of social connections. As tools like Snapchat introduce functions that substitute human-to-human interactions with human-to-AI contact, social isolation could become more pervasive.

Thus, while new AI chatbots and technologies offer 24/7 service and extended reach, we must be cognizant of potential negative impacts on population health and wellbeing (Greenfield & Bhavnani, 2023). With the increasing digitization of human interactions, we risk losing the richness, diversity, and depth inherent in interpersonal relationships. This erosion of human-to-human interactions could increase social isolation and polarization, inadvertently contributing to a decline in population health.

Perpetuating Misinformation

While AI-powered tools offer the potential for improved communication, they also risk facilitating and amplifying misinformation (World Health Organization, 2023). The emergence of AI-driven deepfakes now permit the creation and manipulation of media content, introducing doubt about authenticity when applied to trusted health figures and organizations (Angelis, et al., 2023). At present, social media companies can curate and prioritize content, and their advertisement-based business models can neglect accuracy in favour of sensationalism. The COVID-19 pandemic demonstrated how social media could facilitate the spread of

misinformation, leading to more transmission, vaccine hesitancy, and loss of trust in population health (Suarez-Lledo & Alvarez-Galvez, 2021). Without adequate regulatory intervention, incentive alignment, and digital literacy, the commercial factors that shape content curation in social media could fuel further misinformation with the assistance of AI.

III. List of Recommendations

1. **Expand AIDA's definition of harm to include group harms against a collective**
2. **Expand definition of biased output to include discrimination beyond grounds set out in the *Human Right Code* and remove "without justification" from the definition of biased output contained in section 5(1).**
3. **Expand scope of AIDA to include public sector actors.**
4. **Expanding harm mitigation measures of AI-systems beyond "high-impact systems".**
5. **Expand definition of "persons responsible" to include those involved in training and testing AI systems.**
6. **Mandate a health equity impact assessment by an independent third party to evaluate likelihood of harm and biased output. The assessment must be made publicly available.**
7. **Ensure the advisory committee to the Minister is made up of a diverse membership that reflects a range of communities and perspectives.**
8. **Mandate periodic Parliamentary Review of AIDA**
9. **Enable the Office of the Artificial Intelligence and Data Commissioner to monitor the translational implications and population health impacts of AI-systems.**
10. **Refer AIDA back to the Committee for additional revisions, with a focus on expanding the Act's scope to guarantee accountability and ensuring inclusion of essential details around harm mitigation.**

IV. Recommendations to Committee

Recommendation #1: Expand AIDA's definition of harm to include group harms against a collective

Expanding AIDA's definition of harm to encompass group harms against collectives is a critical step in addressing the complex population health implications of AI systems. A significant concern with the use of AI systems is the risk of systemic bias, which has been repeatedly identified as a key issue by industry and civil organizations alike. AI presents a promising tool to support population health and wellbeing through enhanced data-driven approaches; however, these systems can also conceal unchecked systemic biases that may worsen existing disparities. The embedded systemic bias in AI systems poses a substantial risk, with the potential to intensify prejudices and perpetuate patterns of discrimination and harm when these systems are implemented.

Currently, AIDA's framework primarily addresses harm in terms of individual bodily integrity, property, and economic loss. However, this approach neglects the collective dimension of harm that groups can experience through the adoption of AI systems, such as intensifying racial discrimination, widening economic disparities, and reinforcing societal prejudices. Recognizing and addressing collective harms within AIDA's definition of harm is critical, not only from a legal and ethical standpoint but also for aligning AI development with equity and social justice principles. Recognizing and addressing these collective harms within AIDA's definition of harm is not just a matter of legal and ethical urgency but a crucial step in ensuring that AI development aligns with the principles of equity and social justice, safeguarding the rights and well-being of all individuals and groups, particularly those most vulnerable to exploitation and discrimination.

Recommendation #2: Expand definition of biased output to include discrimination beyond grounds set out in the Human Right Code and remove "without justification" from the definition of biased output contained in section 5(1).

Given the rapid evolution of AI's societal impacts and its training on diverse data sets, we propose broadening the definition of biased output in to encompass currently unaddressed forms of discrimination. A broader definition is crucial for AIDA to effectively guard against a spectrum of discriminatory practices, such as those linked to proxy attributes like postal codes or educational backgrounds, which are outside the current *Human Rights Code* scope. This expansion is vital to prevent adverse impacts on individuals and communities, prevent perpetuating existing societal divides, and ensure AI aligns with Canada's vision of a diverse and inclusive society.

Moreover, the phrase "without justification" in AIDA's section 5(1) implies there could be acceptable circumstances for biased outputs, contradicting the principle of equality that anti-discrimination laws seek to uphold. The inclusion of the phrase also runs contrary to the goal of equity and justice, as it could potentially justify discriminatory practices under the guise of operational efficiency or other factors. Removing "without justification" from the definition would unequivocally affirm that all forms of biased outputs are unacceptable. This amendment

would strengthen the commitment to fairness in AI design and operation, ensuring alignment with human rights principles and Canada's dedication to an equitable society where technology serves to enhance, rather than erode, the principles of equal opportunity and non-discrimination.

Recommendation #3: Expand scope of AIDA to include public sector actors.

AIDA currently does not apply to government institutions as they are defined in the Privacy Act, nor does it apply to products under the control of national security agencies such as the Minister of National Defense or the Director of the Canadian Security Intelligence Service. This exclusion is concerning from an equality and equity perspective. We echo the recommendations made by other organizations to the Committee to broaden AIDA's framework to encompass government institutions.

The integrity and fairness of AI systems used in public sectors like healthcare are paramount, given their broad reach and impact on our communities. Without adequate oversight, there is a risk that AI applications within the health sector, as well as other public services, may cause harm. Expanding AIDA to include public sector entities ensures that AI-driven public sector initiatives adhere to stringent ethical standards. This expansion will protect against biases and reinforce our commitment to health equity. Adopting a holistic approach to AI governance in the public sector will spur innovations that can serve the public good.

Recommendation #4: Expanding harm mitigation measures of AI-systems beyond "high-impact systems".

Currently, AIDA mandates harm mitigation measures solely for high-impact systems, a classification approach that could inadvertently create a regulatory blind spot. By allowing only certain systems to fall under the purview of regulation, it may incentivize developers and users of AI to downplay the potential impact of their systems to avoid stringent oversight. We have already seen this happen with Health Canada's risk-based regulatory framework for medical devices; AI-enabled medical devices have applied for and been granted Class 1 status, the lowest level of risk corresponding with the lowest level of oversight and reporting.

The reality is that AI systems can manifest biases in unexpected ways, often in ways that developers do not foresee and depending on the context they're used in. This has been evidenced by numerous recorded instances of AI bias with significant implications for human rights and substantive equality (Heikkila, 2022) (Dastin, 2018). Therefore, unless the term "high impact" is so broadly defined that it encompasses nearly any AI system with social relevance, the current approach may fail to address the core objective of the legislation: to effectively mitigate harm and prevent discriminatory bias.

To avoid such shortcomings, it is essential to consider a more inclusive strategy for regulatory oversight. This could involve expanding harm mitigation requirements to all AI systems, ensuring that each system is scrutinized for potential harm in its specific context of use. The nature of some applications (e.g. health care) inherently entail elevated levels of risk. The reach of an

application could also be considered. Changes to a platform serving millions of users will carry higher population risks than smaller startups. However, in considering reach, the demographic groups affected must be considered in addition to population sizes. For example, when defining reach, we must consider who the systems affect the most, irrespective of group size. This can help mitigate the risk of propagating systemic inequities existing today. This could also help to promote a competitive AI ecosystem, mitigating the commercial determinants of health, and promoting equity, without stifling Canadian innovation. Alternatively, the responsibility for assessing the impact of AI systems could be placed with an independent third-party assessor, ensuring that the evaluation is impartial and thorough. By implementing these strategies, AIDA would be better positioned to safeguard individuals, groups, and communities from the unintended consequences of AI, thereby aligning more closely with its central aim of mitigating harm and discriminatory bias across the full spectrum of AI applications.

Recommendation #5: Expand definition of "persons responsible" to explicitly include those involved in training and testing AI systems.

AI systems are inherently susceptible to embedded systemic biases, which often originate from their formative phases. As such it's crucial to ensure that those who train and develop AI systems are held accountable for the outcomes. Data-driven bias, a prevalent form of embedded systemic bias, typically arises when AI systems are trained on datasets that fail to represent the diversity of the real world or contain prejudiced information. This can lead to AI outputs that perpetuate existing societal inequities. Algorithmic bias, which refers to systematic biases that stem from the design and decision-making processes within the AI itself, is more elusive. By broadening the accountability to explicitly include professionals who train and test AI systems, there is an opportunity to address these biases proactively, ensuring that AI systems are subjected to rigorous and fair training protocols that reflect the environments they will operate in and the harms they may cause. This comprehensive approach to responsibility not only enhances the fairness and reliability of AI systems but also serves as a proactive measure in safeguarding against the deep-seated biases that AI, without such oversight, could otherwise exacerbate.

Recommendation #6: Mandate a health equity impact assessment by an independent third party to evaluate likelihood of harm and biased output. The assessment must be made publicly available.

Mandating a health equity impact assessment by an independent third party to evaluate the likelihood of harm and biased output from AI systems is a critical measure to ensure proactive harm mitigation. The process of detecting and mitigating the harm caused by AI is often time-consuming and complex. Many individuals may not even realize they have been affected by AI-induced harm unless they possess the capacity and resources to investigate and pursue redress. Moreover, seeking redress can be daunting or impossible for many due to fear of reprisal, lack of resources, or insufficient knowledge to navigate legal channels. Consequently, it's imperative that potential harms from AI systems are identified and addressed well before implementation.

While AIDA aims to mitigate harm via Section 8, it currently lacks explicit requirements for conducting risk assessments at critical stages of AI system development and deployment. Health equity impact assessments, such as the method defined by the Canadian Public Health Association¹, are a practical tool to identify unintended potential health impacts, particularly for vulnerable or marginalized populations. By mandating health equity impact assessments, the government can integrate these considerations into the mandatory assessment that must occur before the deployment of a new AI system. This approach not only provides a systematic method to identify and mitigate potential risks but also reinforces the need for transparency and accountability. Additionally, requiring that the results of such equity and privacy audits are made public helps to ensure that the public is aware of the risks associated with an AI system. This revision would not only help to ensure compliance and transparency but also help to foster a culture of responsibility in AI development and deployment.

Recommendation #7: Ensure the advisory committee to the Minister is made up of a diverse membership that reflects a range of communities and perspectives.

The robust regulatory framework that is meant to follow and build upon AIDA is of paramount importance given the intricate details and provisions it will need to provide. As AIDA outlines, the advisory committee to the Minister will play a pivotal role in providing advice on matters related to the development of a follow up regulatory framework. Ensuring this committee is composed of a diverse membership reflecting a range of communities and perspectives is not just beneficial but essential for the comprehensive development and oversight of AIDA and its regulatory companion. Large technology companies should not be permitted to dominate the conversation given demonstrated challenges with striking a balance between commercial interests and public welfare. Meaningful, in-depth consultation with a wide array of stakeholders, including public health organizations, representatives from communities impacted by AI systems, and not solely industry beneficiaries, is crucial to fostering a more inclusive environment in the regulatory development phase. A varied representation will facilitate a richer understanding of potential harms and foster the creation of effective, nuanced mitigation strategies within the regulatory framework. By prioritizing diversity and inclusivity in the advisory committee, the resulting regulations and the framework at large will be better positioned to protect and represent the interests of all individuals and communities affected by AI.

Recommendation #8: Mandate periodic Parliamentary Review of AIDA.

Mandating a periodic Parliamentary Review of the AIDA is an essential step toward aligning it with the principles of good AI governance, ensuring that the Act not only remains relevant and up to date with scientific advances but also takes into account the fast-changing landscape of AI systems and their use in society. Good AI governance requires that AI systems are bias-free, subject to rigorous audit and evaluation, and equally applicable across all sectors utilizing AI. It emphasizes the need for AI systems to be reliable, their operations and decisions understandable and transparent, and holds those who create and deploy AI systems accountable. Moreover, it

¹ <https://www.cpha.ca/policy-statement-health-equity-impact-assessment>

prioritizes the protection of privacy and insists on a governance framework that emerges from meaningful engagement with all stakeholders, especially those who might bear the brunt of AI's adverse impacts. In a rapidly evolving AI landscape, periodic Parliamentary reviews of AIDA would ensure that the legislation stays updated and reflective of these governance principles.

Recommendation #9. Enable the Office of the Artificial Intelligence and Data Commissioner to monitor the translational implications and population health impacts of AI-systems.

While AI holds enormous potential, algorithms have demonstrated deficiencies through biases and hallucinations. As the pace of progress accelerates, it will be critical to identify issues within systems before they become embedded. The implementation of safeguards to protect the public interest necessitates deliberate and continuous monitoring, extending beyond the scope of individual systems. While AIDA outlines requirements around the measurement and assessment of high-impact AI systems, this must be extended to include monitoring of real-world performance and population-level impacts. Regular assessments should occur, focusing not only on the technical performance of tools but also on the societal and ethical consequences of AI deployments.

At the individual system level, monitoring should be mandated as a prerequisite for the continuous maintenance of licensure. Frameworks must go beyond the technical performance of a tool to capture translational aspects – their impact on workflows and activities in the real world (Reddy, et al., 2021). These should cover stages including the development, deployment, integration, and adoption of AI systems. Ethical dimensions are of particular importance and should include privacy, non-maleficence and explainability.

Equally significant is the establishment of population-level indicators, which are designed to ensure that innovations serve the public interest, and that public welfare remains paramount. Indicators should be developed through the collaborative efforts of a diverse group of stakeholders in order to ensure a comprehensive perspective. It is imperative that part of this monitoring focus specifically on the impacts experienced by equity-deserving groups, ensuring evaluations are unable to mask poor performance in population subgroups. Monitoring of the digital determinants of health – including access to tools, digital literacy, and community infrastructure – can also contribute to promoting the welfare of the public in both the development and implementation of new tools.

Recommendation #10: Refer AIDA back to the Committee for additional revisions, with a focus on expanding the Act's scope to guarantee accountability and ensuring inclusion of essential details around harm mitigation.

While this document outlines some of our recommended revisions to the current legislation, our most significant feedback is that AIDA lacks the scope to fully protect the health of our populations. While this Act may be a reasonable starting point, additional regulatory measures will be needed to adequately address the multifaceted challenges of AI in a manner that upholds ethical standards and protects public welfare. As previously highlighted in recommendation 7, in

the process to revise the Act, the Committee should consult a body with a diverse, representative membership, empowered to apply an equity lens and provide avenues to strengthen the Act's harm mitigation measures.

V. About PHABC

The Public Health Association of BC is a voluntary, non-profit, non-government, member driven organization that provides leadership to promote health, well-being and social equity. We fulfill our mission through advocacy, collaboration, education and research and do this work through the spectrum of public health practice; including prevention, promotion, protection and policy.

VI. Bibliography

- Angelis, L. D., Baglivo, F., Arzilli, G., Privitera, G. P., Ferragina, P., Tozzi, A. E., & Rizz, C. (2023). ChatGPT and the rise of large language models: the new AI-driven infodemic threat in public health. *Frontiers in Public Health*, 11. doi:<https://doi.org/10.3389/fpubh.2023.1166120>
- Artificial Intelligence for Health Task Force. (2020). *Building a Learning Health System for Canadians*. Canadian Institute for Advanced Research.
- Canadian Public Health Association. (2017). *Public Health: A Conceptual Framework*. Canadian Public Health Association.
- Dastin, J. (2018, October 10). *Amazon scraps secret AI recruiting tool that showed bias against women*. From Reuters: <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- European Union Agency for Fundamental Rights. (2022). *Bias in Algorithms: Artificial Intelligence and Discrimination*. Vienna.
- Greenfield, D., & Bhavnani, S. (2023). Social media: generative AI could harm mental health. *Nature*. doi:<https://doi.org/10.1038/d41586-023-01693-8>
- Guo, L. N., Lee, M. S., Kassamali, B., Mita, C., & Nambudiri, V. E. (2022). Bias in, bias out: Underreporting and underrepresentation of diverse skin types in machine learning research for skin cancer detection-A scoping review. *Journal of the American Academy of Dermatology*, 157-159. doi:10.1016/j.jaad.2021.06.884
- Heikkila, M. (2022, March 29). *Dutch scandal serves as a warning for Europe over risks of using algorithms*. From Politico: <https://www.politico.eu/article/dutch-scandal-serves-as-a-warning-for-europe-over-risks-of-using-algorithms/>
- Hold-Lunstad, J. (2022). Social Connection as a Public Health Issue: The Evidence and Systemic Framework for Prioritizing the "Social" in Social Determinants of Health. *Annual Reviews of Public Health*, 193-213. doi:<https://doi.org/10.1146/annurev-publhealth-052020-110732>
- Kickbusch, L., & Holly, L. (2023). Addressing the digital determinants of health: health promotion must lead the charge. *Health Promotion International*. From <https://doi.org/10.1093/heapro/daad059>
- Klein, N. (2023, May 8). *AI machines aren't 'hallucinating'. But their makers are*. From The Guardian: <https://www.theguardian.com/commentisfree/2023/may/08/ai-machines-hallucinating-naomi-klein>
- National Collaborating Centre for Determinants of Health. (2013). *Public Health Roles for Improving Health Equity*.
- Norori, N., Hu, Q., Aellen, F., & Tzovara, A. (2021, October 8). Addressing bias in big data and AI: A call for open science. *CellPress*. doi:<https://doi.org/10.1016/j.patter.2021.100347>
- Pandya, A., & Lodha, P. (2021). Social Connectedness, Excessive Screen Time During COVID-19 and Mental Health: A Review of Current Evidence. *Frontiers in Human Dynamics*, 3. doi:<https://doi.org/10.3389/fhumd.2021.684137>
- Public Health Agency of Canada. (2022). *Pan-Canadian Health Data Strategy: Toward a world-class health data system*.

- Reddy, S., Rogers, W., Makinen, V.-P., Coiera, E., Brown, P., Wenzel, M., . . . Kelly, B. (2021). Evaluation framework to guide implementation of AI systems into healthcare settings. *BMJ Health & Care Informatics*. From <https://pubmed.ncbi.nlm.nih.gov/34642177/>
- Strubell, E., Ganesh, A., & McCallum, A. (2019, June 5). Energy and Policy Considerations for Deep Learning in NLP. *ArXiv*. From <https://arxiv.org/pdf/1906.02243.pdf>
- Suarez-Lledo, V., & Alvarez-Galvez, J. (2021). Prevalence of Health Misinformation on Social Media: Systematic Review. *Journal of Medical Internet Research*, 23(1). doi:10.2196/17187
- World Health Organization. (2023, May 16). *WHO calls for safe and ethical AI for health*. From World Health Organization: <https://www.who.int/news/item/16-05-2023-who-calls-for-safe-and-ethical-ai-for-health>