

Closing the Human Rights Gap in AI Governance



Author

Philip Dawson (Element AI)

Contributors

Grace Abuhamad (Element AI)

Sherif Elsayed-Ali (Element AI)

Marc-Étienne Ouimette (Element AI)

Executive Summary

As different approaches to governing artificial intelligence (AI) have struggled to build public trust, a number of scholars, international organizations and civil society advocates have put forward an alternative model: the international human rights law framework. Universal in scope and benefiting from global legitimacy and state adherence, its proponents argue that a human rights approach to governing AI—with its emphasis on law, rights, accountability and remedy—enjoys a clear value proposition.

In October 2019, Element AI partnered with the Mozilla Foundation and The Rockefeller Foundation to convene a workshop on the human rights approach to AI governance to determine what concrete actions could be taken in the short term to help ensure that respect for human rights is embedded into the design, development, and deployment of AI systems. Global experts from the fields of human rights, law, ethics, public policy, and technology participated. This report provides a summary of the workshop discussions and includes a list of recommendations that came out of the meeting.

The report recommends that governments should adopt a phased-approach to making human rights due diligence and human rights impact assessments a regulatory requirement in the public and private sectors, beginning with the development of model frameworks and sector-specific codes of conduct. The report recommends that industrial policy adopt a human rights approach, for instance through the creation of tailored direct spending programs to help ensure that the design and technological foundations of rights-respecting AI, such as transparency, explainability and accountability, are firmly established in key sectors.

The report also examines the potentially transformative role that a group of investors could play in shaping a new ecosystem of technology companies. Finally, the report recommends that governments implement a dedicated capacity building effort to accelerate understanding of how the existing legislative and regulatory framework can be applied to ensure respect for human rights, and identify potential gaps where adjustments may be necessary. This could be accomplished through the creation of an independent Centre of Expertise on AI, which could assume a range of new functions as a source of policy expertise, capacity building and oversight across government departments, regulatory agencies, industry, civil society, and international organizations.

Table of Contents

Executive Summary	2
Table of Contents	3
Introduction	4
Embedding Human Rights into AI Governance	6
Human Rights Due Diligence	6
Human Rights by Design	10
Data Governance Design	12
Investing in a Responsible AI Fund	13
Equipping Governments to Govern AI	14
Centre of Expertise	15
Governance Toolkit for Rights-Respecting AI	16
What Investors Can Do	16
What Companies Can Do	16
What Governments Can Do	17
Acknowledgements	18

Introduction

Artificial intelligence is expected to generate important social and economic gains, from transforming enterprise productivity to advancing the [Sustainable Development Goals](#). At the same time, recurring revelations of problematic impacts of the use of AI -- such as in the criminal justice system, predictive policing, public social benefits systems, targeted surveillance, advertisements and disinformation -- highlight the extent to which the misuse of AI can threaten universal rights. The rights include privacy, equality and non-discrimination, freedom of assembly, freedom of expression, freedom of opinion, and, in some cases, even democracy.

Over the last few years, a growing number of initiatives¹ have made progress on AI policymaking, and yet few have addressed how AI governance could be adapted to ensure respect for human rights and freedoms, despite a growing body of research and proposals from known scholars, civil society and international organizations.² This is likely to change, however, as international organizations such as the Council of Europe begin to examine the feasibility of a legal framework for artificial intelligence based on human rights standards, democracy and the rule of law; the new President of the European Commission prepares to propose legislation to address the human and ethical implications of AI during her first 100 days in office; and the government of New Zealand launches a collaboration with the World Economic Forum to consider new approaches to regulating AI.

¹ Perhaps the first such efforts were the national AI strategies, which focused mainly on growing the AI sector responsibly by increasing investment in research and training; standing out among an overwhelming number of proposals, the Montreal Declaration for the Responsible Development of Artificial Intelligence was developed through extensive civic consultation to establish an ethical framework for AI that now boasts the signatures of over 75 organizations; technical standards development organizations such as the IEEE and ISO launched projects to guide the ethical and/or trustworthy design of AI systems; the European Commission's High-level Expert Group on AI later developed the Ethics Guidelines for Trustworthy AI as well as Policy and Investment Recommendations for Trustworthy AI; the World Economic Forum and the Partnership on AI have led a series of projects aiming to guide the design, development and deployment of ethical or responsible AI; and, following the development of its G20-endorsed Principles on Artificial Intelligence, the Organisation for Economic Co-operation and Development (OECD) is now preparing to launch an AI Policy Observatory to help countries "encourage, nurture and monitor the responsible development of trustworthy artificial intelligence (AI) systems for the benefit of society".

² Latonero, M. [Governing Artificial Intelligence: Upholding Human Rights & Dignity](#), Data & Society, October 10, 2018; Access Now, [Human Rights in the Age of Artificial Intelligence](#), November 8, 2018; McGregor, L., Murray, D., & Ng, V. (2019). [International Human Rights Law as a Framework for Algorithmic Accountability](#). *International and Comparative Law Quarterly*, 68(2), 309-343; Donahoe, E., Metzger, M. [Artificial Intelligence and Human Rights](#). *Journal of Democracy*, vol. 30 no. 2, 2019, p. 115-126; Council of Europe, [Responsibility and AI: A study of the implications of advanced digital technologies \(including AI systems\) for the concept of responsibility within a human rights framework](#), October, 9 2019; Australian Human Rights Commission, World Economic Forum. White Paper: Artificial Intelligence: governance and leadership, 2019.

Clearly, much work remains to be done, including addressing the underlying business models of large platforms.³ This report contributes to existing efforts that aim to close the “human rights gap” in AI governance, and contains a series of recommendations that could help embed respect for human rights into the design, development, and deployment of AI systems: first, by embedding human rights due diligence into corporate governance; second by incentivizing research and development into human rights by design; third, by developing innovative models of data governance that are capable of incorporating the core elements of the human rights framework, such as data trusts; and through the creation of national, independent Centres of Expertise on AI, which could assume a range of new functions as a source of policy expertise, capacity building and capacity across government policymaking departments, regulatory agencies, industry, civil society, the United Nations and other multilateral stakeholders.



³ Surveillance Giants: How the Business Model of Google and Facebook Threatens Human Rights. November 2019. Amnesty International. Retrieved at: <https://www.amnesty.org/download/Documents/POL3014042019ENGLISH.PDF>



Embedding Human Rights into AI Governance

Respect for human rights must be built into AI systems from the start. Conducting human rights due diligence and impact or risk assessments is emerging as a basic requirement that all companies and governments using AI will need to undertake in order to comply with key legal requirements and to maintain a social licence with citizens and customers. Human rights law principles also help highlight the design and technological imperatives of rights-respecting AI, calling for increased investments in research and development on the transparency, explainability and accountability of AI systems. Finally, innovative approaches to data governance that are capable of embedding the features of the human rights framework (including respect for rights, accountability and remedy) should be explored. Data trusts are one potential example of such an approach.

Human Rights Due Diligence

Governance tools like **human rights due diligence (HRDD)** and **human rights impact assessments (HRIA)** can help governments and businesses address the human rights impact of artificial intelligence. Although not currently required by law, they have become an established best practice in a variety of economic sectors, notably, within the extractive, agricultural, and consumer goods industries. For companies operating in the global garment and retail industries, HRDD and HRIA are a fixture of corporate governance policy. With the adoption of the UN [Guiding Principles on Business and Human Rights](#) in 2011, there is now a globally recognized and authoritative framework for conducting ongoing human rights due diligence.⁴

⁴Corporate human rights due diligence – identifying and leveraging emerging practice, retrieved from <https://www.ohchr.org/EN/Issues/Business/Pages/CorporateHRDueDiligence.aspx>: “Human rights due diligence is a way for enterprises to proactively manage potential and actual adverse human rights impacts with which they are involved. It involves four core components: (a) Identifying and assessing actual or potential adverse human rights impacts that the enterprise may cause or contribute to through its own activities, or which may be directly linked to its operations, products or services by its business relationships; (b) Integrating findings from impact assessments across relevant company processes and taking appropriate action according to its involvement in the impact; (c) Tracking the effectiveness of measures and processes to address adverse human rights impacts in order to know if they are working; (d) Communicating on how impacts are being addressed and showing stakeholders – in particular affected stakeholders – that there are adequate policies and processes in place.”

In light of the clear human rights risks posed by AI systems, **governments should begin a phased-approach to making HRDD and HRIA a regulatory requirement.** First, governments should support the development of model frameworks for HRDD and HRIA that identify what factors must be included in these types of assessments (e.g. adequate consequences if identified risks of adverse human rights impacts are not duly mitigated and addressed).⁵ The creation of model frameworks for HRDD and HRIA could help mitigate the risk of them becoming rubber-stamping exercises,⁶ and would be necessary to facilitate uptake and implementation. Second, because of the context specific nature of these types of assessments, model HRDD and HRIA would also need to be adapted to the particular sectors in which AI is being deployed.

Third, while the clearest case for requiring HRDD and HRIA is in the context of the procurement of AI systems by public institutions, including all levels of government and the courts,⁷ a risk-based approach should guide their mandatory use in the private sector.⁸ In particular, mandatory implementation in the private sector should focus on high-risk applications of AI or instances where human rights violations, such as discrimination, could have a significant effect on a person's financial interests, personal health or well-being, or where minors are affected. The use of AI systems in financial services, the education system, human resources and healthcare should be considered as priority areas; government regulators should also make it a requirement for social media platforms to base their terms of service or community standards on universal human rights principles and engage in HRDD or HRIA, given the increasing role they play in the adjudication of universal rights such as privacy,

⁵ Recommendation: Unboxing artificial intelligence: 10 steps to protect human rights, By the Council of Europe Commissioner for Human Rights, May 2019, retrieved from: <https://rm.coe.int/unboxing-artificial-intelligence-10-steps-to-protect-human-rights-reco/1680946e64>

⁶ For an example of a possible risk assessment model, see Figure 2: Criticality pyramid and risk-adapted regulatory system for the use of algorithmic systems in the "Opinion of the Data Ethics Commission", The Data Ethics Commission of the Federal Government of Germany (October 2019), retrieved at https://datenethikkommission.de/wp-content/uploads/191023_DEK_Kurzfassung_en_bf.pdf

⁷ Momentum is already building for this approach. In Canada, the Directive on Automated Decision-Making Systems and accompanying Algorithmic Impact Assessment developed by the federal government is an example of an HRIA that has been tailored to the administrative context, where rights to due process, procedural fairness and appeal make the transparency, explainability and accountability of AI systems a legal requirement. In November 2019, Waterfront Toronto, the public entity leading the development of the City of Toronto's waterfront area, released a request for proposals to perform a HRIA of the Master Innovation and Development Plan presented by Sidewalk Labs, the affiliate of Alphabet Inc. that plans to perform part of the work. The tender states that the HRIA should be based off existing international best practices and informed by the United Nations Guiding Principles on Business and Human Rights. The decision to seek a preliminary HRIA of the company's innovation plan follows an extended period of controversy, notably, in relation to Sidewalk's data governance proposals.

⁸ Legislative requirements for corporate due diligence already exist in other sectors to control human rights risks, including in the United Kingdom's Modern Slavery Act 2015, the European Union Timber Regulation and in the French Duty of Vigilance Law. Privacy impact assessments, required in comprehensive data protection legislation, are other helpful examples.

discrimination and the freedom of opinion and expression through their content moderation responsibilities.⁹ While risk assessments can be used to help identify, control or minimize risks,¹⁰ they may also help identify instances in which the risks to human rights and people's well-being are simply too high, and lead to a decision to place a moratorium on deployment.¹¹

The imposition of due diligence requirements or impact assessments is particularly appropriate to emerging technologies like AI, which can be a moving target for regulators. **Rather than making the technology itself the object of regulation, this approach helps ensure companies develop a systematic and auditable way of controlling and minimizing risks, while allowing space for innovation.** Risk assessments have an excellent track record of helping regulators integrate highly innovative technologies and approaches in a variety of sectors. The success of the aviation industry in using risk assessments to consistently reduce the risk of fatalities over its history, and more recently in integrating new entrants such as small remotely piloted aircraft systems (commonly known as drones) into national airspace, serves as but one example.

The **requirement to perform HRDD and HRIA grounded in international and domestic human rights law should be supplemented by reporting requirements and audits performed by an appropriate government entity.** The European Union Directive 2014/95/EU on non-financial reporting, which requires the reporting of human rights impact assessments by large undertakings, could serve as a helpful example in this regard. Reporting requirements and audits of HRDD and HRIA could introduce greater transparency into how such assessments are conducted, incentivizing more meaningful accountability for technology companies.

As more technology companies perform HRIA and HRDD, **standards should be developed by industry and professional bodies to further entrench respect for human rights in the design, development and deployment of AI.** The IEEE included "human rights" as its first "General Principle" in guidance it produced *Ethically Aligned Design*, making explicit reference to the international human rights framework and the relevance of the UN Guiding Principles to the design, development and deployment of AI. Standard development organizations such as ISO, and national chapters, could adopt

⁹ Kaye, D. Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression, 2018 thematic report to the Human Rights Council on content regulation, retrieved from <https://documents-dds-ny.un.org/doc/UNDOC/GEN/G18/096/72/PDF/G1809672.pdf?OpenElement>

¹⁰ Vermeulen, M. Online Content: To Regulate or Not to Regulate is that the Question? October 2019. APC Issue Paper. Retrieved at: <https://www.apc.org/sites/default/files/OnlineContentToRegulateOrNotToRegulate.pdf>

¹¹ Molnar, 2019, "Technologies on the Margins: The Human Rights Impacts of AI in Migration Management." Cambridge International Law Journal, Volume 8, Issue 2, 2019.

a human rights-based approach to the work currently being conducted in ISO/IEC JTC 1/SC 42 - Artificial intelligence, in particular, with respect to the development of an explainability framework for AI.

Beyond helping to promote legal compliance, **the business case for embedding HRDD grounded in the Universal Declaration of Human Rights through the UN Guiding Principles is also attractive.** For instance, the UN Guiding Principles can help provide:

- A proactive risk management framework that merges legal, human rights and ethical considerations that captures a comprehensive portrait of potential harm (including risk of liability, as many universal human rights are incorporated into domestic legislation and regulation, e.g. privacy and discrimination).
- A realistic forecast of the true costs of the design, development and deployment of AI, by highlighting the technological imperatives of rights-respecting AI (e.g. human-centered design, transparency, explainability and accountability).
- Access to a growing pool of resources and guidance on due diligence best practices (including international jurisprudence, national case law, best practices in HRDD and HRIAs).
- An international framework that can be scaled across global operations; and
- In some cases, companies that have satisfied due diligence requirements may be able to raise these as a defense in the event of liability claims.

This past year, the United Nations Office of the High Commissioner for Human Rights launched the [B-Tech Project](#) to help governments, companies and other stakeholders understand how the UN Guiding Principles can be tailored and applied to the context of AI. The project focuses on the following four areas: 1) addressing human rights in business models; 2) human rights due diligence and end use; 3) accountability and remedy; 4) exploring regulatory and policy responses to human rights challenges linked to digital technologies.¹²

¹² UN Human Rights, UN Human Rights Business and Human Rights in Technology Project (B-Tech): Applying the UN Guiding Principles on Business and Human Rights to digital technologies Draft Scoping Paper for Consultation July, 30 2019. Retrieved from: <https://www.ohchr.org/EN/Issues/Business/Pages/B-TechProject.aspx>.



Human Rights by Design

The **human rights framework highlights that the transparency, explainability and accountability of AI systems is not only a desirable ethical achievement, but a prerequisite to upholding existing legal obligations.** For instance, the lack of transparency of the use and operation of such systems can make it difficult to determine whether a breach of human rights has occurred, pre-empting the ability to seek redress. Similarly, the reliance on a recommendation, decision or insight provided by an AI system that is not explainable is at odds with human rights and administrative law, which in many cases require that an individual be provided with reasons for a decision made against them; the transparency and explainability of AI systems is therefore a legal prerequisite to their accountability.

Governments should **clarify that a general right to an explanation already exists when an AI system informs a decision that has a significant effect on a person's rights, financial interests, personal health or well-being.**¹³ Because what is required by way of an explanation will vary according to the decision-making circumstances and the objectives of the system in question, approaches to the explainability of AI systems will be met with unique technological feasibility challenges. Accordingly product developers should consider human rights by design principles such as explainability at the earliest ideation stages to the later deployment stages.

¹³ Edward Santow, Briefing on Australian Human Rights Commission's Forthcoming Human Rights and Technology Discussion Paper (17 October 2019).



To ensure that the technological foundations of rights-respecting AI are established as a matter of priority, policymakers and practitioners in government and business should:

- **Disincentivize irresponsible technology deployment through regulation, but also incentivize research and development for work on responsible AI (e.g., explainability) in the private sector through tailored direct spending programs, or other financial incentives;**
- **Ensure that procurement of AI systems by public institutions place a strong emphasis on technological and design solutions,** in particular, to promote transparency, explainability and accountability of AI systems; and
- **Create new discretionary spending programs to help expand existing research initiatives rooted in privacy by design, to human rights by design.** Invest in action-oriented research and tools to assist human rights by design policies and practices.

Governments should include the importance of adopting policies and programs that incentivize the development of rights-respecting AI in their national AI strategies to ensure they remain a priority.

Data Governance Design

Fiduciary models of data governance such as data trusts¹⁴ have received a growing amount of consideration recently, notably, as private law mechanisms capable of introducing civic representation and accountability into the use of data and AI. Data trusts also have the potential to incorporate certain elements of the human rights framework into data and AI governance, namely, through the trustee's fiduciary duty to protect data subjects' rights and by providing a collective action mechanism that enables accountability in case of harm.

While a growing body of research has focused on the legal and governance foundations of data trusts, a high level of automation will be needed to help data trustees manage an increasing amount of requests for data access in a manner that ensures informed consent and concords with the purposes of the trust. In addition, accountability mechanisms will need to be built into the technical architecture of a data trust to ensure that access to a remedy is readily available at scale. Governments can encourage the development of data governance models that enable the respect for human rights by:

- **Creating an enabling regulatory framework for data trusts, including incorporating them into comprehensive data protection legislation;**
- **Supporting the development of the technological underpinnings of data trusts with direct spending programs; and**
- **Creating pilot programs requiring participation from industry, academia and civil society to encourage the testing of models like data trusts.**

¹⁴ Data trusts result from the application of the common law trust to the governance of data or data rights. Trusts begin with an asset, or rights in an asset, that a "settlor" places into a trust. A trust charter stipulates the purpose and terms of the trust, which exists to benefit a group of people, known as the "beneficiary". In more basic terms, a data trust creates a legal way to manage data rights for a purpose that is valuable to a beneficiary. In a data trust, individuals would be empowered to pool the rights they hold over their personal data into the legal framework of a trust. A trustee is appointed with a fiduciary obligation to manage the trust's assets in accordance with the trust charter and the interests of its beneficiaries. The trustee is accountable to the beneficiaries for the management of the trust, and has a responsibility to take legal action to protect their rights. For more, please see: Delacroix, S. & Lawrence, N. (2019). Bottom-up data Trusts: disturbing the 'one size fits all' approach to data governance. *International Data Privacy Law*. <https://doi.org/10.1093/idpl/ipy014>; Dawson, P. & Abuhamad, G. (2019), *Towards Data Governance that Empower the Public*, Global Information Society Watch 2019: Artificial intelligence: Human rights, social justice and development (forthcoming); McDonald, S. (2019). Reclaiming Data Trusts. *Centre for International Governance Innovation*. <https://www.cigionline.org/articles/reclaiming-data-trusts>; Rahaak, A. Data Trusts: Why, What and How, 2019, <https://medium.com/@anoukruhaak/data-trusts-why-what-and-how-a8b53b53d34>.



Investing in a Responsible AI Fund

In the absence of government regulation, investment firms can play a disproportionate role in conditioning the AI market by injecting billions of dollars into technology companies that aim for speed, scale, and profit. When board members are concerned with short-term metrics, commitments to ongoing due diligence or the fundamental, applied and design research programs necessary for building rights-respecting AI can be difficult for management teams to justify.

Fortunately, the tide appears to be changing, as organizations and companies that champion **“digital human rights” are increasingly cited as the “next frontier for fund groups”**.¹⁵ Investors can play an important role in developing the AI ecosystem by empowering the companies they invest in to adopt the research budgets, corporate governance structures and timelines for market returns that recognize the imperatives of responsible, rights-respecting AI.

And if individual investors can impact the market one company at a time, what if they combined to fill the **need for a new Responsible AI Fund capable of incentivizing and supporting the long term development needs associated with rights-respecting AI**. A global coalition of venture capital firms, foundations, governments and other institutional investors - including social impact and sovereign wealth funds - could raise enough “patient” capital to disrupt the status quo and support the growth of responsible AI companies.

¹⁵ November 10, 2019. Digital human rights are next frontier for fund groups. Financial Times. retrieved from <https://www.ft.com/content/0866d79f-cd48-42d4-b21c-453f964d2fb0>.



Equipping Governments to Govern AI

National AI strategies have committed substantial funds towards AI research and talent development to grow the AI sector; a capacity building effort is now needed to empower governments with the personnel and institutional arrangements needed to ensure that AI is designed, developed and deployed in a manner that protects human rights.

Government policymakers and regulators must be equipped with personnel capable of monitoring and assessing AI's domain-specific risks, especially regarding the rights, health and financial interests of those affected. As such, **a dedicated capacity building effort within government is necessary to accelerate understanding of how the existing legislative and regulatory framework can be applied to ensure respect for human rights**, and identify potential gaps where adjustments may be necessary.

For instance, departments of justice must understand the risk of algorithmic bias in predictive policing and judicial-decision making; officials administering the public benefits system require expertise on algorithmic bias and discrimination; immigration, tax and revenue departments must ensure that the integration of AI respects human rights and principles by administrative law (which depend on the transparency, explainability and accountability of AI systems). The protection of internet users from the risk of online hate speech or the spread of disinformation requires expertise on internet governance, content moderation techniques and the freedom of opinion and expression. Policy makers contemplating the reform of privacy legislation should understand how AI-driven methods of analysis on social media can be used to infer and generate sensitive information about people that they have neither provided nor confirmed, turning the consent-based approach to privacy protection on its head.



Centre of Expertise

Capacity-building within government could be accomplished through the creation of an independent **Centre of Expertise on AI to monitor, assess, report on, and provide advice to government and industry about risk management in AI**, or through a decentralized effort modelled after Digital Service programs instituted in the United Kingdom, United States, and Canada.¹⁶ The Centre of Expertise on AI should be an agile source of policy expertise, education and capacity building. The Centre of Expertise could:

- **Support the development of Model Frameworks for HRDD and HRIA**, including sector-specific codes of conduct in priority sectors.
- **Undertake research and publish reports on the social impact and human rights risks posed by AI systems.** Scholars, civil society and international organizations have published research outlining AI's human rights risks; the Centre of Expertise should build on this work to provide society with a comprehensive overview of the human rights risks triggered by AI systems, and empower individuals with knowledge of the recourses available to them (e.g., under anti-discrimination law or human rights codes).
- **Ensure that departments and ministries understand how AI may affect their respective regulatory roles, notably through the human rights lens.**
- **Issue notices or advisories on AI applications that pose high risks to human rights**, in coordination with regulators. Similar to the function performed by aviation, tax or securities authorities, this type of action could help raise social awareness with respect to high risk uses of AI systems, while conditioning the market to the risk of liability.
- As the national brain trust of AI governance, **contribute research and domestic best practices to international and multilateral efforts aimed at guiding the harmonization of responsible AI, grounded in human rights**, such as the Canada-France led Global Partnership on AI.

¹⁶ Edward Santow, Briefing on Australian Human Rights Commission's Forthcoming Human Rights and Technology Discussion Paper (17 October 2019); to institutionalize a risk-management approach to governing AI in Canada, Dan Munro has previously recommended that the Government of Canada create two new institutions: an AI risk governance council and an algorithm impact assessment agency. Please see Munro, D. Governing AI: Navigating Risks, Rewards and Uncertainty, Public Policy Forum, January 11, 2019. Retrieved from: <https://ppforum.ca/publications/governing-ai/>.



Governance Toolkit for Rights-Respecting AI

AI governance requires a collaborative multi-stakeholder approach. One of the goals of the workshop was to build on the research and thought leadership undertaken by academic and research communities, civil society, and international organizations by developing a list of regulatory, governance and policy recommendations that companies, governments and investors can act on in the near term.

What Investors Can Do

- Support companies that adopt research budgets, corporate governance structures and timelines for market returns that recognize the imperatives of responsible, rights-respecting AI;
- Fund research and advocacy efforts designed to empower the public with knowledge of AI systems and their risks, in particular to human rights; and
- Explore the possibility of establishing a new Responsible AI Fund capable of incentivizing and supporting the long term development needs associated with rights-respecting AI.

What Companies Can Do

- Commit to conducting HRDD and HRIA throughout the AI lifecycle;
- Support and contribute to the UN Office of the High Commissioner for Human Rights' B-Tech project, to ensure the UN Guiding Principles are properly adapted to the context of AI;

-
- Prioritize research on the design and technological imperatives of rights-respecting AI systems, with a focus on transparency, explainability and accountability; and
 - Collaborate with partners from academic communities, civil society, international organizations and governments to help them understand the risks associated with AI systems and work together to devise appropriate governance mechanisms and safeguards.

What Governments Can Do

- Support a phased approach to requiring HRDD and HRIA in the public and private sectors, beginning with developing model frameworks and sector-specific codes of conduct that may be audited;
- Establish a new, independent Centre of Expertise on AI (as more fully described above);
- Disincentivize irresponsible technology deployment through regulation, but also incentivize research on human rights by design in the private sector, with an emphasis on transparency, explainability and accountability, through tailored direct spending programs, or other financial incentives;
- Incentivize research and development of the technological underpinnings of data trusts through new spending and pilot programs;
- Support the Canada-France led Global Partnership on AI, which could serve as a forum for international coordination of research and best practices developed by national Centres of Expertise on AI; and,
- Partner with the Office of the UN High Commissioner for Human Rights to host national consultations for the B-Tech Project, to support the application of the UN Guiding Principles to the context of AI.



Acknowledgements

This report follows an Element AI workshop on AI Governance and Human Rights that took place in October 2019 at The Rockefeller Foundation's Bellagio Center. This multi-stakeholder convening would not have been possible without generous support from the Mozilla Foundation and The Rockefeller Foundation.

Element AI would like to express its appreciation to the following list of experts for their participation in the workshop, and for their comments on drafts of this report:

Aaina Agarwal, AI & Human Rights Consultant

J. Bob Alotta, Mozilla Foundation

Charlotte Altenhoener-Dion, Council of Europe;

Steven Crown, Microsoft Corporation;

Sylvie Delacroix, Alan Turing Institute;

Eileen Donahoe, Stanford Global Digital Policy Incubator;

Peggy Hicks, Office of the United Nations High Commissioner for Human Rights;

Fanny Hidvegi, Access Now;

Mark Latonero, Data & Society;

Roya Pakzad, Taraaz;

Chloe Poynton, Article One Advisors;

Tobias Rees, Berggruen Institute;

Edward Santow, Australian Human Rights Commission;

Mark Surman, Mozilla Foundation

Evan Tachovsky, The Rockefeller Foundation

Richard Whitt, Glia Foundation;

Residents of The Rockefeller Foundation's [AI thematic month](#); with special thanks to **Rumman Chowdhury**, Accenture, and **Sherry Wong**, Icarus Salon, for their participation.

Element AI would also like to thank the following experts who provided thoughtful feedback on preparations for the workshop or on versions of this report: **Tim Engelhardt**, **Kyle Matthews**, **Petra Molnar**, **Dan Munro**, **Nicole Rigillo**, **Philippe-André Rodriguez**, and **Trooper Sanders**.

Element AI's effort in AI Governance and Human Rights builds on the company's participation in the development of European High-level Expert Group's [Ethics Guidelines for Trustworthy AI](#), the Organization for Economic Cooperation and Development's [Principles on Artificial Intelligence](#), and reflects the company's ongoing commitment to developing an enabling environment for AI deployment that is socially beneficial and sustainable in the long term. Element AI is excited to be continuing this work with the Mozilla Foundation as a [new partner and collaborator](#) in building responsible AI.

Montreal

6650 Saint-Urbain
Suite #500
Montreal, QC, H2S 3G9
Canada
+1 (514) 379-3568

Toronto

296 Richmond St. W
Suite #100
Toronto, ON, M5V 1X2
Canada

London

Vox Studios
Suite W-111
45 Durham, Vauxhall
London SE11 5JH
United Kingdom

Seoul

Dreamplus Gangnam
311 Gangnam-daero
Seocho-gu, Seoul
Republic of Korea

Singapore

168 Robinson Road
#24-01 Capital Tower
068912, Singapore

CONTACT US
elementai.com/contact

