



PROJECT MUSE®

Artificial Intelligence and Human Rights

Eileen Donahoe, Megan MacDuffee Metzger

Journal of Democracy, Volume 30, Number 2, April 2019, pp. 115-126 (Article)

Published by Johns Hopkins University Press

DOI: <https://doi.org/10.1353/jod.2019.0029>



➔ *For additional information about this article*

<https://muse.jhu.edu/article/721650>

ARTIFICIAL INTELLIGENCE AND HUMAN RIGHTS

Eileen Donahoe and Megan MacDuffee Metzger

Eileen Donahoe, executive director of the Global Digital Policy Incubator and adjunct professor at Stanford University's Center on Democracy, Development, and the Rule of Law, is former U.S. ambassador to the UN Human Rights Council in Geneva. **Megan MacDuffee Metzger** is research scholar and associate director for research at the Global Digital Policy Incubator.

In democratic societies, concern about the consequences of our growing reliance upon artificial intelligence (AI) is rising. The term AI, coined by John McCarthy in 1956, is elusive in its precise meaning but today broadly refers to machines that can go beyond their explicit programming by making choices in ways that mirror human reasoning. In other words, AI automates decisions that people used to make.¹

While AI promises many benefits, there are also risks associated with the swift advancement and adoption of the technology. Perhaps the darkest concerns relate to misuse of AI by authoritarian regimes. Even in free societies, however, and even when the intended application is for clearly good purposes, there is significant potential for unintended harms such as reduced privacy, lost accountability, and embedded bias. In digitally connected democracies, talk of what could go wrong with AI now touches on everything from massive job loss caused by automation to machines that make discriminatory hiring decisions, and even to threats posed by “killer robots.” These concerns have darkened public attitudes and made this a key moment to either build or destroy public trust in AI.

How did we get to this point? In the connected half of the world, the shift to the “data-driven” society has been quick and quiet—so quick and quiet that we have barely begun to come to grips with what our growing reliance on machine-made decisions in so many areas of life will mean for human agency, democratic accountability, and the enjoyment of human rights.

Many governments have been formulating national AI strategies to keep from being left behind by the AI revolution, but few have been grap-

pling seriously with AI's implications for their accountability or their duty to protect citizens' rights. Some private companies (Google and Microsoft among them) have set forth their own ethical principles for developing and applying AI and have embraced accountability for how their AI-based products affect users and customers.² Individual technologists, academics, and civil society actors are working to develop guidelines to help ensure that AI is used for human benefit. These initiatives can improve awareness of AI's risks, but none offers a comprehensive framework for addressing these risks, nor can any claim anything approaching global buy-in.

A shared global framework is needed to ensure that AI is developed and applied in ways that respect human dignity, democratic accountability, and the bedrock principles of free societies. We argue that the existing universal human-rights framework is well suited to serve this function and already has wide global legitimacy. Advocating the broad utilization of this existing framework and articulating how to apply it in the context of digital societies are the important next steps toward governing AI.

Societal and Ethical Concerns About AI

To understand how the existing human-rights framework can help us to address the human implications of AI, it is useful first to delineate several different types of potential harm and layers of societal concern. The grimmest worries about AI—that machines will come to control society or that autonomous weapons systems will turn against their human creators—have captured imaginations, but in fact relate to a type of technology that does not yet exist and may never exist. This is artificial general intelligence (AGI), which by definition would match or exceed the reasoning powers of humans. Experts debate whether AGI is achievable and if so in what timeframe. To date, there is still disagreement about what human intelligence *is*, let alone whether machines might someday come to outstrip it.

A second level of society-wide debate centers around questions of how to weigh whether or when various applications of AI are ethical, who should make such judgments, and on what basis. Some private firms, for example, have expressed deep misgivings about whether they should sell certain AI applications—especially facial-recognition technologies—to governments. Facial recognition has benign and even positive applications.³ Many people use it as an added layer of security for their phones. Yet it hardly requires an out-of-control imagination to worry about facial recognition as a potent tool for state surveillance and violation of civil liberties. Modern cities already bristle with cameras, and their use in the constant, ubiquitous monitoring and surveillance of law-abiding citizens is already happening in places such as the People's Republic of China. Microsoft and Google have faced employee revolts

over company plans to sell certain AI products to governments. This in turn has led Google executives to decline renewal of a U.S. government contract and Microsoft executives to call for government regulation of facial recognition. At this juncture, many in the private sector now acknowledge that deeper civic engagement is needed when it comes to governance of AI and its applications.⁴

Even when AI serves worthy ends, a third debate centers on *unintended* negative effects. For example, biases embedded in data used to train AI can become reified and magnified under the guise of objectivity.⁵ Digitization, which has pushed us toward greater reliance on AI, also presents deep challenges to privacy and creates new systemic security vulnerabilities. How do we protect privacy and security in digitized societies where everyone's personal data—whether they fully consent or not—is used to train AI? While automation promises to yield impressive economic efficiencies, at least in the short term, it may force people out of work and exacerbate inequality. These effects make AI suspect even when applied for the sake of good ends.

At a more foundational level, AI also is testing our assumptions about human agency. When individuals use private-sector digital platforms to access information, AI systems that are intended to optimize results in light of the user's personal preferences may have the paradoxical effect of undermining the exercise of free choice by determining the kinds of information to which people are exposed. With respect to reliance on AI in public-sector governance, we must address questions about whether human beings will be held accountable for machine decisions that have an impact on people's rights. Until now, human societies have been structured and governed with human beings as the focal point. We must be alert to the risk that AI systems may not be geared toward enhancing human dignity. As AI advances, the status of human beings as the focal point of AI-reliant governance decisions cannot be assumed—it will have to be ensured.

A range of practical problems that AI might cause for democratic accountability, due process, and the rule of law has been manifest in the recent adoption of AI by several U.S. jurisdictions that apply it in matters of criminal sentencing, parole, eligibility for social services, and employment decisions. AI-based systems—which are inherently opaque—can undermine transparency and fairness in public-sector decision making, especially when even the people who are formally responsible for certain decisions do not understand the basis on which “the algorithm” is making these choices. This is especially troubling if existing biases are baked into the data that feed the decisions.

Public accountability for government decisions is crucial to democracy: Officials must be held responsible for decisions that violate citizens' rights. If governments are going to rely upon AI, they must develop processes to evaluate how machine results affect people's rights, and be ready to provide timely, effective remedies in cases where machine-

made decisions turn out to have been wrong. Simply put, as governments shift further toward reliance on AI, we need mechanisms to ensure democratic accountability.

A Human-Centered Ethics for AI

In response to rising public concerns about AI, many stakeholders have begun talking about new ethical frameworks to mitigate its risks and ensure its beneficial application. These initiatives include efforts by individual companies such as Google and Microsoft to develop their own ethical principles for AI, but also collaborative efforts between private companies and other stakeholders and organizations, including the Partnership on AI (*partnershiponai.org*), the Asilomar Principles,⁶ and Open AI (*openai.com*). Civil society and the academy also have ramped up efforts to develop new ethical frameworks for AI. The organization Fairness, Accuracy, and Transparency in Machine Learning (*fatml.org*) has brought together a consortium of different stakeholders to uphold those standards in machine learning and to encourage governments and the private sector to commit to take steps that will guard against bias in AI-based decisions. On the academic front, Stanford University has launched its Human-Centered AI Institute (*hai.stanford.edu*), and the Institute of Electrical and Electronics Engineers (IEEE) has brought scholars together with technologists and civil society activists to launch the Global Initiative on Ethics of Autonomous and Intelligent Systems.⁷

These are all valuable efforts to address various risks posed by AI. While each effort has hold of part of the problem, none grasps the full spectrum of risks to humans or the whole range of governance challenges. Perhaps more importantly, none of these initiatives can claim to have achieved a broad global consensus or buy-in across stakeholder groups. We need a comprehensive, global governance framework to address the full spectrum of AI challenges.

While private firms and other, nonprofit groups deserve credit for clarifying their own commitments and publishing ethical principles for AI-related work, the influence of these efforts beyond company or consortia walls remains uncertain. All lack normative sway over entities that took no part in drafting them. Ethics statements may guide the entities that commit to them, but they do not establish a broad governance framework under which all can operate.

If we are going to have any chance of effectually facing the broad range of challenges that AI poses for society, we will need something that goes beyond the particular and takes a more systematic and comprehensive approach. We will also need a framework that can claim global buy-in and that addresses the roles and responsibilities of both government and the private sector when it comes to accountability for the impact of AI-based decisions. Fortunately, the foundation for this more

comprehensive and global approach already exists in a well-elaborated form. It is the universal human-rights framework that was first laid down under the auspices of the United Nations in 1948.

The Universal Declaration of Human Rights (UDHR), drafted in the aftermath of the Second World War and adopted by vote of the UN General Assembly meeting in Paris on 10 December 1948, established a universally applicable set of individual rights and government commitments to protect people. It took a global crisis of unspeakable proportions to establish these governance principles. Human rights as conceived in the Declaration inhere in every human person, without regard to geography or regime. Following the UDHR came a series of legally binding international treaties that explicate the wide range of civil, political, economic, social, and cultural rights envisioned in the founding Declaration. Together, these documents—known as the International Bill of Human Rights—create a body of international human-rights law that establishes the obligations of governments to protect (and not violate) the rights of citizens and other people within those governments' respective territories and jurisdictions. Under the human-rights framework, governments have the duty to protect citizens from violations and infringements of their rights by other governments and nonstate actors, including the private sector.

A human-rights–based approach to governance of AI can accomplish what the newly emerging ethical frameworks seek to accomplish. Four particular features of the human-rights framework make this possible: 1) It puts the human person at the center of any assessment of AI and makes impact on humans the focal point of governance; 2) it covers the wide range of pressing concerns that AI raises, both procedural and substantive; 3) it outlines the roles and duties of both governments and the private sector in protecting and respecting human rights; and 4) it has the geopolitical advantage of resting on a broad global consensus, shared by many countries around the world and understood to be universally applicable. Together, these features make a strong case to support use of this existing framework to guide the governance of AI.

First, the human-rights framework starts with a conception of human dignity and asserts the centrality of the human person as the focal point of governance and society. This commitment is the critical first step toward ensuring that AI is developed in ways that support human beings and human societies, rather than in ways that might harm them. If used to govern AI, the human-rights framework will require assessments of AI technologies based on what they mean for people and their inherent rights. This starting point is essential to ensure sufficient consideration of the impact on humans, especially in cases where machines are increasingly taking on roles that people previously held.

Second, this international body of human-rights law, through its broad spectrum of both substantive and procedural rights, speaks directly to the most pressing societal concerns about AI. More specifically,

the following sections of the UDHR directly address critical societal concerns about the potential impacts of AI:

Article 2: The right to equal protection and nondiscrimination speaks to concerns about avoiding bias in data and ensuring fairness in machine-based decisions relied upon by governments, especially when these decisions impact citizens' rights.

Article 3: The right to life and personal security can form the foundation for addressing concerns about autonomous weapons that move beyond human control, as well as the risk of a literal arms race to develop such weapons.

Articles 8 through 11: Together, the right to effective remedy for violations and infringements of rights and the right to due process of law provide a basis for addressing concerns about transparency, fairness, and accountability in cases where AI systems impact people's rights.

Article 12: The right to privacy addresses the fundamental concern about loss of privacy in data-driven societies and the need to protect personally identifiable data.

Article 19: The right to freedom of expression, which includes the right to hold opinions and to seek, share, and access information, speaks to challenges associated with the increasing role played by algorithms in moderating content and curating what information is presented to users of digital platforms.

Articles 20 and 21: The rights to association and peaceful assembly, as well as the right to democratic participation in government, are becoming increasingly important as lenses through which to assess the impact of AI systems on human agency, the organization of civic groups, political participation, and the integrity of democratic processes around the globe.

Articles 23 and 25: The rights to work and to enjoy an adequate standard of living—rights whose fulfillment, unlike the fulfillment of first-generation negative rights such as freedom of expression, does not rest solely with governments—can guide governance decisions around displacement of human workers by AI.

Third, the human-rights framework establishes the roles and responsibilities of both governments and the private sector in protecting and respecting human rights and in remedying violations of them. The original framework established the obligation of governments to protect the rights of citizens or others under their jurisdiction in ways that none of the newly emerging ethical frameworks can match. In addition, the framework requires states to ensure effective remedies against rights violations by non-state actors, which is key to accountability. The UN Guiding Principles on Business and Human Rights (UNGPs), adopted by the UN Human Rights Council in 2011, more substantially articulate the role and responsibili-

ties of private-sector businesses in protecting human rights.⁸ Within the UNGP framework, the general legal obligation to protect human rights remains with states, while private firms have responsibilities to respect and protect human rights (and to remedy violations of them) when the firm's own products, services, and operations are involved. Given the central role that private-sector technology companies play in developing and applying AI, this normative development was significant.

The UNGPs are especially well suited to AI-driven societies where private-sector AI products and services affect so many areas of life, often in opaque ways. The UNGPs establish due-diligence standards that private companies can follow in order to identify and mitigate risks to human rights. Under the UNGPs, private companies should embrace ongoing responsibility to perform human-rights–impact assessments. The guidelines also establish expectations that firms will consult stakeholders and publish reports of their own due-diligence processes.

In sum, the International Bill of Human Rights and the UNGPs provide a fairly comprehensive framework for governance of AI, as well as for assessing the impacts on people and society of decisions made by AI. This combination of global norms, legal agreements, and process expectations addresses one of the most pressing questions around AI and digital technology in general: What are the relative responsibilities of governments as distinguished from tech companies when it comes to the effects that new AI technologies cause? No other proposed AI-governance framework defines the rights of individuals, the obligations of governments, and the responsibilities of private companies with the comprehensiveness of the human-rights framework.

Finally, although interpreted and implemented in vastly different ways around the world, the existing universal framework enjoys a level of geopolitical recognition and status under international law that no newly emergent ethical framework can match. This is a crucial advantage. The human-rights framework rests on decades of global consensus. It has been embedded in national constitutions and applied by governments around the world in national regulations. It is the product of global negotiation reflecting the work of many stakeholders, and establishes a universally applicable set of norms and commitments.⁹ Countries that do not comply with these norms risk censure from the international community. This does not mean that any states fully comply with these standards, or that all states fully embrace these principles as guiding norms. But it is to say that human-rights standards enjoy a high level of legitimacy.

Indeed, at a rhetorical level even states that do not genuinely support or try to comply with human-rights principles often attempt to characterize their behavior in ways that reflect human-rights language and standards. Pragmatically speaking, moreover, it is unlikely, given current global political trends, that were a new global governance framework for AI to be negotiated, it would encompass the full spectrum of com-

mitments to human dignity that the existing human-rights framework contains. While authoritarian governments might like to shift global support away from human-rights norms in the digital context, democratically aligned governments should vehemently resist such a shift.

Simply put, in contrast to ad hoc ethical frameworks generated by isolated groups of concerned parties, universal human-rights principles have served as a globally recognized set of governance norms for the past seventy years. Although interpreted and implemented in different ways around the world, this framework has nonetheless been critical to forging a global consensus about the spectrum of rights that people everywhere should enjoy. Is the framework, which dates back decades before the web, still relevant in an age of fast-moving digital technologies? In 2012, the UN Human Rights Council in effect asserted that it was when the Council adopted, by consensus, the first UN resolution on the promotion and protection of human rights on the internet. The Council affirmed the foundational concept that “the same rights that people have offline must also be protected online.”¹⁰

Since then, in service to this basic commitment, helpful analysis explaining how existing universal human-rights norms can apply in newly digitized settings has begun to emerge.¹¹ Articulating how to apply these human-rights principles to AI will take more work. But beginning with global standards to which states have agreed brings enormous legitimacy, if not always perfect implementation. A human-rights–based approach provides the best opportunity to protect against the worst risks of AI while allowing people to benefit from its many opportunities.

Much remains to be done, but existing international human-rights law can provide governments, private companies, civil society, and citizens at large with a shared and indeed global approach to preventing (or at least mitigating) the harms that might flow from AI. As machines grow more capable, new AI applications can and should be evaluated with reference to their impact on the broad spectrum of rights that inhere in each human person. In sum, there is no need to reinvent the wheel or to start with a blank page: Existing international human-rights law and principles are the best resource we have for dealing with the wide range of social and ethical concerns that AI has already raised and will continue to raise.

Governing AI Through a Human-Rights Lens

We are at a critical juncture when it comes both to how we govern AI and how we use AI in governance. As technology advances, laws and norms struggle to keep pace. Governments and the international community have been forced to play catch-up as tech-driven crises erupt. The need for rules and regulations to govern AI’s development and use is urgent. That said, the question should *not* be whether we need a new ethical foundation for the governance of AI—we already have such a

foundation in the existing human-rights framework. Our task, rather, is to discern how to apply this suitable framework to new digital technologies.

Shifting from questions about what principles should guide our assessments of AI to questions about how to implement a human-rights–based approach to AI will require creative thinking across disciplines. Computer scientists will need to work together with human-rights theorists, philosophers, international lawyers, and policy experts. Human-rights considerations must make themselves felt in the design, application, and evaluation of AI products, as well as in all government procurement and deployment of AI. This can increase the likelihood that the massive potential of AI will benefit societies, without causing harm to the human beings who compose them.

It is critical that governments and other international stakeholders commit themselves more explicitly to support a human-rights–based approach to the governance of AI, and that they help to explain why such an approach has the greatest potential to address the full range of societal concerns that AI raises. If we continue with an ad hoc approach to the ethics of AI, or if we allow AI to go ungoverned, the outcome will not be optimal for anyone. Without a global set of standards, countries will vie to achieve innovations, but without necessarily considering the negative externalities these might cause. Thus it is imperative to call explicitly for an approach based squarely on human rights. If most members of the international community can coalesce around a human-rights–based approach to the governance of AI that is consistent with their existing commitments, this will help to bring along more reluctant states. It will also ward off efforts by authoritarian-leaning states to undermine confidence in the feasibility of protecting human rights in the digital context.

The next step is to articulate more clearly how to implement human-rights principles in all sectors of an AI-driven world. So far, civil society groups and international organizations have taken the lead in this task. For example, the UN’s special rapporteur on freedom of expression, David Kaye, has begun applying human-rights law to evaluate AI’s impact on free expression when AI is used to moderate and curate content on digital platforms.¹² The Council of Europe’s human-rights commissioner, Dunja Mijatović, is a leading advocate of the human-rights framework for AI governance.¹³ In the United Kingdom, the Human Rights, Big Data, and Technology Project at the University of Essex has asked Parliament to adopt such a framework to shape government regulation of AI.¹⁴ Mark Latonero of Data and Society has convened academics and civil society activists to discuss new ways of using the human-rights framework to address AI’s societal impacts.¹⁵ Access Now, a leading civil society group, has written several reports advocating the use of existing international human-rights principles to regulate AI.¹⁶ Article One Advisors works with companies to develop their own processes for assessing the human-rights impact of AI-based products and services.

Several key practical ideas have emerged from these efforts. The first is the fundamental need for transparency in both governmental and business uses of decision-making algorithms. David Kaye has urged private digital platforms such as Amazon, Facebook, and Twitter to be far more open with their users about how AI curates online information. This includes helping users to understand how algorithms influence which information they see and which information appears most often. He has also recommended that private platforms allow users to opt out of AI curation, something that Twitter has recently implemented.¹⁷

Another new idea that is now drawing the attention of product developers is the concept of “human rights by design.”¹⁸ This rests on a willingness to ask the question “What will this technology mean for the human rights of the people who use it?” This work of reflection must take place *as* technologies are being developed—by means of human-rights-impact assessment processes, for instance—not *after* they have already been adopted. In order to fulfill their obligation to protect citizens, public authorities also need to step up their vigilance regarding what private companies are doing with AI. At a more mundane level, governments should be required to think carefully about their own AI-procurement decisions, as well as the consequences of deploying AI products in the administration of public services that bear on citizens’ rights.

As we articulate more clearly how human-rights standards can be applied to AI in private-sector practices and public policy alike, the international community as a whole, as well as states individually, must also develop mechanisms of accountability to ensure that both governments and the private sector are complying with these standards. There must be clear consequences for failures to comply, as well as remedies when violations occur. While it is fine for private firms to devise their own ethical frameworks, global technology companies would be wise to emphasize the existing universal human-rights framework and the UNGPs in assessing what their products and services will mean for people. Providing remedies is very important, but anticipating damage and heading it off are better still.¹⁹

We also need to teach young technologists about existing human-rights norms so that those who seek to build ethical AI applications need never feel that they are operating in a vacuum. In this regard, it is exciting to note that some technologists have begun to talk about designing human-rights commitments into AI products and services.²⁰ “Human rights by design” is such a promising concept that schools and universities should make its study a priority.

Finally, we need more dialogue across sectors and more cross-disciplinary education. That includes technical education for policy makers, human-rights education for technologists, and both human-rights and AI education for all students. The challenges associated with governance *of* AI and governance *by* AI cannot be tackled by

any one group; only broad cross-sector engagement can do the trick. Putting human rights first offers the best hope for protecting people from potential harms and for building societies that are enriched and strengthened by what AI technology has to offer.

Societal challenges brought by AI demand new approaches to governance, but these challenges do not require new normative principles. The human-rights framework that we already have is well suited to the global digital environment. As applications of AI proliferate, so must practical ways of bringing human-rights standards to bear. Our urgent task is to figure out how to protect and realize human rights in our new AI-driven world.

NOTES

1. For an introduction, see Chris Smith et al., “The History of Artificial Intelligence,” University of Washington, December 2006, <https://courses.cs.washington.edu/courses/csep590/06au/projects/history-ai.pdf>. On how AI is affecting rights, see Levin Kim et al., “Artificial Intelligence and Human Rights: Opportunities and Risks,” Berkman Klein Center for Internet and Society, Harvard University, 25 September 2018, <https://cyber.harvard.edu/publication/2018/artificial-intelligence-human-rights>.

2. “Artificial Intelligence at Google: Our Principles,” Google, <https://ai.google/principles/>; “Microsoft AI Principles,” www.microsoft.com/en-us/ai/our-approach-to-ai.

3. “Delhi: Facial Recognition System Helps Trace 3,000 Missing Children in 4 Days,” *Times of India*, 22 April 2018; Jeannine Mjoseh, “Facial Recognition Software Helps Diagnose Rare Genetic Disease,” National Human Genome Research Institute, 23 March 2017.

4. Natasha Singer, “Microsoft Urges Congress to Regulate Use of Facial Recognition,” *New York Times*, 13 July 2018; Daisuke Wakabayashi and Scott Shane, “Google Will Not Renew Pentagon Contract that Upsets Employees,” *New York Times*, 1 June 2018; Jake Laperruque and Andrea Peterson, “Amazon Pushes ICE to Buy Its Face Recognition Surveillance Tech,” *Daily Beast*, 23 October 2018, www.thedailybeast.com/amazon-pushes-ice-to-buy-its-face-recognition-surveillance-tech.

5. Julia Angwin et al., “Machine Bias: There’s Software Used Across the Country to Predict Future Criminals. And It’s Biased Against Blacks,” *Pro Publica*, 23 May 2016, www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing; Jason Tashea, “Courts Are Using AI to Sentence Criminals: That Must Stop Now,” *Wired*, 17 April 2017, www.wired.com/2017/04/courts-using-ai-sentence-criminals-must-stop-now.

6. “Asilomar AI Principles,” Future of Life Institute, n.d., <https://futureoflife.org/ai-principles/?cn-reloaded=1>.

7. See <https://standards.ieee.org/industry-connections/ec/autonomous-systems.html>.

8. “Guiding Principles on Business and Human Rights: Implementing the United Nations ‘Protect, Respect, and Remedy’ Framework,” 16 June 2011, www.ohchr.org/Documents/Publications/GuidingPrinciplesBusinessHR_EN.pdf.

9. UN General Assembly, “International Covenant on Civil and Political Rights,” 16 December 1966, United Nations, Treaty Series 999 (16 December 1966): 171, www.refworld.org/docid/3ae6b3aa0.html; UN General Assembly, “International Covenant on Economic, Social and Cultural Rights,” 16 December 1966, United Nations, Treaty Series 993 (16 December 1966): 3, www.refworld.org/docid/3ae6b36c0.html.

10. UN Human Rights Council, Resolution 20/8, “The Promotion, Protection and Enjoyment of Human Rights on the Internet,” 16 July 2012.

11. See, for example, “Necessary and Proportionate: International Principles on the Application of Human Rights to Communications Surveillance,” May 2014, <https://necessaryandproportionate.org/principles>; “Report of U.N. Special Rapporteur Frank La Rue on the Right to Freedom of Opinion and Expression,” 17 April 2013, www.aclu.org/other/report-un-special-rapporteur-frank-la-rue-right-freedom-opinion-and-expression; “Report of the Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression, David Kaye,” 22 May 2015, www.un.org/en/ga/search/view_doc.asp?symbol=A/HRC/29/32; “The Right to Privacy in the Digital Age,” n.d., www.ohchr.org/en/issues/digitalage/pages/digitalageindex.aspx; Jermyn Brooks and Susan Morgan, “GNI Writes to UN High Commissioner for Human Rights on Privacy in the Digital Age,” 24 February 2014, <https://globalnetworkinitiative.org/gni-writes-to-un-high-commissioner-for-human-rights-on-privacy-in-the-digital-age>.

12. “Report of the Special Rapporteur on the Promotion and Protection of the Right to Freedom of Expression,” 29 August 2018, <https://freedex.org/wp-content/blogs.dir/2015/files/2018/10/AI-and-FOE-GA.pdf>.

13. See “Safeguarding Human Rights in the Era of Artificial Intelligence—Human Rights Comment by Dunja Mijatović, Commissioner for Human Rights,” Council of Europe, 3 July 2018, www.coe.int/en/web/genderequality/-/safeguarding-human-rights-in-the-era-of-artificial-intelligence; and Dunja Mijatović, “We Need to Act Now and Put Human Rights at the Centre of Artificial Intelligence Designs,” Council of Europe, 26 February 2019, www.coe.int/en/web/commissioner/-/we-need-to-act-now-and-put-human-rights-at-the-centre-of-artificial-intelligence-designs.

14. “Written Evidence Submitted by the Human Rights, Big Data and Technology Project (ALG0063),” U.K. Parliament, April 2017, <http://data.parliament.uk/written-evidence/committeeevidence.svc/evidencedocument/science-and-technology-committee/algorithms-in-decisionmaking/written/69117.html>.

15. Mark Latonero, “Governing Artificial Intelligence: Upholding Human Rights and Dignity,” Data and Society, 10 October 2018, <https://datasociety.net/output/governing-artificial-intelligence>.

16. “Human Rights in the Age of Artificial Intelligence,” Access Now, November 2018, www.accessnow.org/cms/assets/uploads/2018/11/AI-and-Human-Rights.pdf; “Mapping Regulatory Proposals for Artificial Intelligence in Europe,” Access Now, November 2018, www.accessnow.org/cms/assets/uploads/2018/11/mapping_regulatory_proposals_for_AI_in_EU.pdf.

17. Will Oremus, “Twitter Has Finally Made It Easy to Set Your Timeline to Reverse-Chronological,” *Slate*, 18 December 2018, <https://slate.com/technology/2018/12/twitter-reverse-chronological-timeline-setting.html>.

18. Jonathon Penney, et al., “Advancing Human-Rights-by-Design in the Dual-Use Technology Industry,” *Columbia Journal of International Affairs*, 20 December 2018, <https://jia.sipa.columbia.edu/advancing-human-rights-design-dual-use-technology-industry>.

19. Microsoft’s human-rights-impact assessments for its AI products and services provide a good model to follow. See www.microsoft.com/en-us/corporate-responsibility/human-rights.

20. Dunstan Allison-Hope, “Human Rights by Design,” 17 February 2017, www.bsr.org/en/our-insights/blog-view/human-rights-by-design.